

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problems Mailbox.**

THIS PAGE BLANK (USPTO)



AH

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 : G06F 7/22, 11/10		A1	(11) International Publication Number: WO 91/13399 (43) International Publication Date: 5 September 1991 (05.09.91)
<p>(21) International Application Number: PCT/US91/01276 (22) International Filing Date: 28 February 1991 (28.02.91)</p> <p>(30) Priority data: 488,749 2 March 1990 (02.03.90) US 505,622 6 April 1990 (06.04.90) US 506,703 6 April 1990 (06.04.90) US 601,482 22 October 1990 (22.10.90) US</p> <p>(71) Applicant: SF2 CORPORATION [US/US]; 140 Kifer Court, Sunnyvale, CA 94086 (US).</p> <p>(72) Inventors: IDLEMAN, Thomas, E. ; 2660 Brady Court, Santa Clara, CA 95051 (US). KOONTZ, Robert, S. ; 140 Walsh Avenue, Atherton, CA 94025 (US). POWERS, David, T. ; 2265 Bayo Claros Circle, Morgan Hill, CA 95037 (US). JAFFE, David, H. ; 551 South Road, Belmont, CA 94002 (US). HENSON, Larry, P. ; 3325 St. Ignatius Place, Santa Clara, CA 95051 (US). GLIDER, Joseph, S. ; 3292 Murray Way, Palo Alto, CA 94303 (US). GAJJAR, Kumar ; 1700 Fan Street, San Jose, CA 95131 (US).</p>			
<p>(74) Agent: HAUGHEY, Paul, C.; Townsend and Townsend, One Market Plaza, 2000 Steuart Tower, San Francisco, CA 94105 (US).</p> <p>(81) Designated States: AT, AT (European patent), AU, BB, BE (European patent), BF (OAPI patent), BG, BJ (OAPI patent), BR, CA, CF (OAPI patent), CG (OAPI patent), CH, CH (European patent), CM (OAPI patent), DE, DE (European patent), DK, DK (European patent), ES, ES (European patent), FI, FR (European patent), GA (OAPI patent), GB, GB (European patent), GR (European patent), HU, IT (European patent), JP, KP, KR, LK, LU, LU (European patent), MC, MG, ML (OAPI patent), MR (OAPI patent), MW, NL, NL (European patent), NO, RO, SD, SE, SE (European patent), SN (OAPI patent), SU, TD (OAPI patent), TG (OAPI patent).</p> <p>Published With international search report.</p>			
<p>(54) Title: DISK ARRAY SYSTEM</p> <p>(57) Abstract</p> <p>A method and apparatus for controlling data flow between a computer (10) and a group of memory devices (18A-18F) arranged in a particular logical configuration. The system includes a group of first level controllers (12A, 12B) and a group of second level controllers (14A, 14B). The first level controllers and the second level controllers work together such that if one of the second level controllers fails, the routing between the first level controllers and the memory devices is switched to a properly functioning second level controller without the need to involve the computer in the rerouting process. The logical configuration of the memory devices remains constant. The invention also includes switching circuitry (16) which permits a functioning second level controller to assume control of a group of memory devices formerly primarily controlled by the failed second level controller. In addition, the invention provides error check and correction (Figure 10) as well as mass storage device configuration circuitry.</p>			
<pre> graph TD COMPUTER[COMPUTER] --> CONTROLLER[CONTROLLER] CONTROLLER --> 12A[TWO-LEVEL DEVICE 12A] CONTROLLER --> 12B[TWO-LEVEL DEVICE 12B] 12A -- MA --> SWITCH[SWITCH] 12B -- MB --> SWITCH SWITCH --> 18A[DISK DRIVE ARRAY] SWITCH --> 18B[DISK DRIVE ARRAY] SWITCH --> 18F[DISK DRIVE ARRAY] </pre>			

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MN	Mongolia
BE	Belgium	GA	Gabon	MR	Mauritania
BP	Burkina Faso	GB	United Kingdom	MW	Malawi
BG	Bulgaria	GN	Guinea	NL	Netherlands
BJ	Benin	GR	Greece	NO	Norway
BR	Brazil	HU	Hungary	PL	Poland
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	SD	Sudan
CG	Congo	KP	Democratic People's Republic of Korea	SE	Sweden
CH	Switzerland	KR	Republic of Korea	SN	Senegal
CI	Côte d'Ivoire	LI	Liechtenstein	SU	Soviet Union
CM	Cameroon	LK	Sri Lanka	TD	Chad
CS	Czechoslovakia	LU	Luxembourg	TG	Togo
DE	Germany	MC	Monaco	US	United States of America
DK	Denmark				

DISK ARRAY SYSTEM

5

CROSS REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part of Serial Nos. 07/505,622, 07/506,703, and 07/488,749, filed April 6, 1990, April 6, 1990, and March 2, 1990, respectively.

10

BACKGROUND OF THE INVENTION

The present invention relates generally to memory storage devices. More particularly, the invention is a method and apparatus for interfacing an external computer to a set of storage devices which are typically disk drives.

Magnetic disk drive memories for use with digital computer systems are known. Although many types of disk drives are known, the present invention will be described as using hard disk drives. However, nothing herein should be taken to limit the invention to that particular embodiment.

Many computer systems use a plurality of disk drive memories to store data. A common known architecture for such systems is shown in Figure 1. Therein, computer 10 is coupled by means of bus 15 to disk array 20. Disk array 20 is comprised of large buffer 22, bus 24, and a plurality of disk drives 30. The disk drives 30 can be operated in various logical configurations. When a group of drives is operated collectively as a logical device, data stored during a write operation can be spread across one or more members of an array. Disk controllers 35 are connected to buffer 22 by bus 24. Each controller 35 is assigned a particular disk drive 30.

Each disk drive within disk drive array 20 is accessed and the data thereon retrieved individually. The disk controller 35 associated with each disk drive 30 controls the input/output operations for the particular disk drive to which it is coupled. Data placed in buffer 22 is available for transmission to computer 10 over bus 15. When the computer transmits data to be written on the disks, controllers 35

receive the data for the individual disk drives from bus 24. In this type of system, disk operations are asynchronous in relationship to each other.

In the case where one of the controllers experiences 5 a failure, the computer must take action to isolate the failed controller and to switch the memory devices formerly under the failed controller's control to a properly functioning other controller. The switching requires the computer to perform a number of operations. First, it must isolate the failed 10 controller. This means that all data flow directed to the failed controller must be redirected to a working controller.

In the system described above, it is necessary for 15 the computer to be involved with rerouting data away from a failed controller. The necessary operations performed by the computer in completing rerouting requires the computer's attention. This places added functions on the computer which may delay other functions which the computer is working on. As a result, the entire system is slowed down.

Another problem associated with disk operations, in 20 particular writing and reading, is an associated probability of error. Procedures and apparatus have been developed which can detect and, in some cases, correct the errors which occur during the reading and writing of the disks. With relation to 25 a generic disk drive, the disk is divided into a plurality of sectors, each sector having the same, predetermined size. Each sector has a particular header field, which gives the sector a unique address, a header field code which allows for the detection of errors in the header field, a data field of variable length and ECC ("Error Correction Code") codes, which 30 allow for the detection and correction of errors in the data.

When a disk is written to, the disk controller reads the header field and the header field code. If the sector is the desired sector and no header field error is detected, the new data is written into the data field and the new data ECC is 35 written into the ECC field.

Read operations are similar in that initially both the header field and header field error code are read. If no header field errors exist, the data and the data correction

codes are read. If no error is detected the data is transmitted to the computer. If errors are detected, the error correction circuitry located within the disk controller tries to correct the error. If this is possible, the corrected data 5 is transmitted. Otherwise, the disk drive's controller signals to the computer or master disk controller that an uncorrectable error has been detected.

In Figure 2 a known disk drive system which has an associated error correction circuit, external to the individual 10 disk controllers, is shown. This system uses a Reed-Solomon error detection code both to detect and correct errors. Reed-Solomon codes are known and the information required to generate them is described in many references. One such reference is Practical Error Correction Design for Engineers, 15 published by Data Systems Technology Corp., Broomfield, Colorado. For purposes of this application, it is necessary to know that the Reed-Solomon code generates redundancy terms, herein called P and Q redundancy terms, which terms are used to detect and correct data errors.

20 In the system shown in Figure 2, ECC 42 unit is coupled to bus 45. The bus is individually coupled to a plurality of data disk drives, numbered here 47, 48, and 49, as well as to the P and Q term disk drives, numbered 51 and 53 through Small Computer Standard Interfaces ("SCSIs") 54 through 25 58. The American National Standard for Information Processing ("ANSI") has promulgated a standard for SCSI which is described in ANSI document number X3.130-1986.

Bus 45 is additionally coupled to large output buffer 22. Buffer 22 is in turn coupled to computer 10. In this 30 system, as blocks of data are read from the individual data disk drives, they are individually and sequentially placed on the bus and simultaneously transmitted both to the large buffer and the ECC unit. The P and Q terms from disk drives 51 and 53 are transmitted to ECC 42 only. The transmission of data and 35 the P and Q terms over bus 45 occurs sequentially. The exact bus width can be any arbitrary size with 8-, 16- and 32-bit wide buses being common.

After a large block of data is assembled in the buffer, the calculations necessary to detect and correct data errors, which use the terms received from the P and Q disk drives, are performed within the ECC unit 42. If errors are 5 detected, the transfer of data to the computer is interrupted and the incorrect data is corrected, if possible.

During write operations, after a block of data is assembled in buffer 22, new P and Q terms are generated within ECC unit 42 and written to the P and Q disk drives at the same 10 time that the data in buffer 22 is written to the data disk drives.

Those disk drive systems which utilize known error correction techniques have several shortcomings. In the systems illustrated in Figures 1 and 2, data transmission is 15 sequential over a single bus with a relatively slow rate of data transfer. Additionally, as the error correction circuitry must wait until a block of data of predefined size is assembled in the buffer before it can detect and correct errors therein, there is an unavoidable delay while such detection and 20 correction takes place.

As stated, the most common form of data transmission in these systems is serial data transmission. Given that the bus has a fixed width, it takes a fixed and relatively large amount of time to build up data in the buffer for transmission 25 either to the disks or computer. If the large, single buffer fails, all the disk drives coupled thereto become unusable. Therefore, a system which has a plurality of disk drives which can increase the rate of data transfer between the computer and the disk drives and more effectively match the data transfer 30 rate to the computer's maximum efficient operating speed is desirable. The system should also be able to conduct this high rate of data transfer while performing all necessary error detection and correction functions and at the same time provide an acceptable level of performance even when individual disk 35 drives fail.

Another failing of prior art systems is that they do not exploit the full range of data organizations that are possible in a system using a group of disk drive arrays. In

other words, a mass storage apparatus made up of a plurality of physical storage devices may be called upon to operate as a logical storage device for two concurrently-running applications having different data storage needs. For example, 5 one application requiring large data transfers (i.e., high bandwidth), and the other requiring high frequency transfers (i.e., high operation rate). A third application may call upon the apparatus to provide both high bandwidth and high operating rate. Known operating techniques for physical device sets do 10 not provide the capability of dynamically configuring a single set of physical storage devices to provide optimal service in response to such varied needs.

It would therefore be desirable to be able to provide 15 a mass storage apparatus, made up of a plurality of physical storage devices, which could flexibly provide both high bandwidth and high operation rate, as necessary, along with high reliability.

SUMMARY OF THE INVENTION

20 The present invention provides a set of small, inexpensive disk drives that appears to an external computer as one or more logical disk drives. The disk drives are arranged in sets. Data is broken up and written across the disk drives in a set, with error detection and correction redundancy data being generated in the process and also being written to a 25 redundancy area. Backup disk drives are provided which can be coupled into one or more of the sets. Multiple control systems for the sets are used, with any one set having a primary control system and another control system which acts as its backup, but primarily controls a separate set. The error 30 correction or redundancy data and the error detection data is generated "on the fly" as data is transferred to the disk drives. When data is read from the disk drives, the error detection data is verified to confirm the integrity of the 35 data. Lost data from a particular disk drive can be regenerated with the use of the redundancy data and a backup drive can be substituted for a failed disk drive.

The present invention provides an arrangement of disk drive controllers, data disk drives and error correction code disk drives, the drives being each individually coupled to a small buffer memory and a circuit for error detection and correction. A first aspect of the present invention is error detection and correction which occurs nearly simultaneously with the transfer of data to and from the disk drives. The multiple buffer memories can then be read from or written to in sequence for transfers on a data bus to the system computer.

5 Additionally, the error correction circuitry can be connected to all of the buffer memory/disk drive data paths through a series of multiplexer circuits called cross-bar ("X-bar") switches. These X-bar switches can be used to decouple failed buffer memories or disk drives from the system.

10 A number of disk drives are operatively interconnected so as to function at a first logical level as one or more logical redundancy groups. A logical redundancy group is a set of disk drives which share redundancy data. The width, depth and redundancy type (e.g., mirrored data or check

15 20 data) of each logical redundancy group, and the location of redundant information therein, are independently configurable to meet desired capacity and reliability requirements. At a second logical level, blocks of mass storage data are grouped into one or more logical data groups. A logical redundancy group may be divided into more than one such data group. The width, depth, addressing sequence and arrangement of data

25 30 blocks in each logical data group are independently configurable to divide the mass data storage apparatus into multiple logical mass storage areas each having potentially different bandwidth and operation rate characteristics.

A third logical level, for interacting with application software of a host computer operating system, is also provided. The application level superimposes logical application units on the data groups to allow data groups, alone or in combination from one or more redundancy groups, to appear to application software as single logical storage units.

35 As data is written to the drives, the error correction circuit, herein called the Array Correction Circuit

("ACC"), calculates P and Q redundancy terms and stores them on two designated P and Q disk drives through the X-bar switches. In contrast to the discussed prior art, the present invention's ACC can detect and correct errors across an entire set of disk
5 drives simultaneously, hence the use of the term "Array Correction Circuit." In the following description, the term ACC will refer only to the circuit which performs the necessary error correction functions. The codes themselves will be referred to as Error Correction Code or "ECC". On subsequent
10 read operations, the ACC may compare the data read with the stored P and Q values to determine if the data is error-free.

The X-bar switches have several internal registers. As data is transmitted to and from the data disk drives, it must go through a X-bar switch. Within the X-bar switch the
15 data can be clocked from one register to the next before going to the buffer or the disk drive. The time it takes to clock the data through the X-bar internal registers is sufficient to allow the ACC to calculate and perform its error correction tasks. During a write operation, this arrangement allows the P
20 and Q values to be generated and written to their designated disk drives at the same time as the data is written to its disk drives, the operations occurring in parallel. In effect the X-bar switches establish a data pipeline of several stages, the plurality of stages effectively providing a time delay circuit.
25

In one preferred embodiment, two ACC units are provided. Both ACCs can be used simultaneously on two operations that access different disk drives or one can be used if the other fails.

The X-bar switch arrangement also provides
30 flexibility in the data paths. Under control of the system controller, a malfunctioning disk drive can be decoupled from the system by reconfiguring the appropriate X-bar switch or switches and the data that was to be stored on the failed disk can be rerouted to another data disk drive. As the system
35 computer is not involved in the detection or correction of data errors, or in reconfiguring the system in the case of failed drives or buffers, these processes are said to be transparent to the system computer.

In a first embodiment of the present invention, a plurality of X-bar switches are coupled to a plurality of disk drives and buffers, each X-bar switch having at least one data path to each buffer and each disk drive. In operation a failure of any buffer or disk drive may be compensated for by rerouting the data flow through a X-bar switch to any operational drive or buffer. In this embodiment full performance can be maintained when disk drives fail.

In another embodiment of the present invention, two ACC circuits are provided. In certain operating modes, such as when all the disk drives are being written to or read from simultaneously, the two ACC circuits are redundant, each ACC acting as a back-up unit to the other. In other modes, such as when data is written to an individual disk drive, the two ACCs work in parallel, the first ACC performing a given action for a portion of the entire set of drives, while the second ACC performs a given action which is not necessarily the same for a remaining portion of the set.

In yet another embodiment, the ACC performs certain self-monitoring check operations using the P and Q redundancy terms to determine if the ACC itself is functioning properly. If these check operations fail, the ACC will indicate its failure to the control system, and it will not be used in any other operations.

In still another embodiment, the ACC unit is coupled to all the disk drives in the set and data being transmitted to or from the disk drives is simultaneously recovered by the ACC. The ACC performs either error detection or error correction upon the transmitted data in parallel with the data transmitted from the buffers and the disk drives.

The present invention provides a speed advantage over the prior art by maximizing the use of parallel paths to the disk drives. Redundancy and thus fault-tolerance is also provided by the described arrangement of the X-bar switches and ACC units.

Another aspect of the present invention is that it switches control of disk drive sets when a particular

controller fails. Switching is performed in a manner transparent to the computer.

The controllers comprise a plurality of first level controllers each connected to the computer. Connected to the

5 other side of the first level controllers is a set of second level controllers. Each first level controller can route data to any one of the second level controllers. Communication buses tie together the second level controllers and the first level controllers can also communicate between themselves. In
10 a preferred embodiment, the system is configured such that the second level controllers are grouped in pairs. This configuration provides each second level controller with a single associated back-up controller. This configuration provides for efficient rerouting procedures for the flow of
15 data to the disk drives. For ease of understanding, the specification will describe the system configured with pairs of second level controllers. Of course, it should be understood that the second level controllers could be configured in groups of three or other groupings.

20 A switching function is implemented to connect each of the second level controllers to a group of disk drives. In the case that a second level controller should fail, the computer need not get involved with the rerouting of data to the disk drives. Instead, the first level controllers and the
25 properly working second level controller can handle the failure without the involvement of the computer. This allows the logical configuration of the disk drives to remain constant from the perspective of the computer despite a change in the physical configuration.

30 There are two levels of severity of failures which can arise in the second level controllers. The first type is a complete failure. In the case of a complete failure, the second level controller stops communicating with the first level controllers and the other second level controller. The
35 first level controllers are informed of the failure by the properly working second level controller or may recognize this failure when trying to route data to the failed second level controller. In either case, the first level controller will

switch data paths from the failed second level controller to the properly functioning second level controller. Once this rerouted path has been established, the properly functioning second level controller issues a command to the malfunctioning 5 second level controller to release control of its disk drives. The properly functioning second level controller then assumes control of these disk drive sets.

The second type of failure is a controlled failure where the failed controller can continue to communicate with 10 the rest of the system. The partner second level controller is informed of the malfunction. The properly functioning second level controller then informs the first level controllers to switch data paths to the functioning second level controller. Next, the malfunctioning second level controller releases its 15 control of the disk drives and the functioning second level controller assumes control. Finally, the properly functioning second level controller checks and, if necessary, corrects data written to the drives by the malfunctioning second level controller.

20 A further aspect of the present invention is a SCSI bus switching function which permits the second level controllers to release and assume control of the disk drives.

For a more complete understanding of the nature and 25 the advantages of the invention, reference should be made to the ensuing detail description taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram illustrating a prior art 30 disk array system;

Fig. 2 is a block diagram illustrating a prior art disk array system with an error check and correction block;

Fig. 3 is a diagram illustrating a preferred embodiment of the overall system of the present invention;

35 Fig. 4 is a diagram showing a more detailed illustration of Fig. 3 including the interconnections of the switches and the disk drives within the disk drive sets;

Fig. 5 is a block diagram of the wiring between the controllers and the switches;

Fig. 6 is a block diagram showing the schematic circuitry of the switching function control circuitry shown in Fig. 5;

Fig. 7 is a recovery state transition diagram illustrating the various possible states of a particular second level controller;

Figs. 8A-8I show the events which take place during the transition between each of the states shown in Fig. 7;

Figure 9 is a block diagram of one preferred embodiment of the X-bar circuitry;

Figure 10 is a block diagram of a preferred embodiment of the error check and correction circuitry;

Figure 11 is a detailed block diagram of the X-bar switches and the ACC shown in Figure 10;

Figures 12a and 12b show the logic operations necessary to calculate the P and Q error detection terms;

Figures 13a and 13b show how the Reed-Solomon codeword is formed and stored in one embodiment of the present invention;

Figures 14a and 14b show the parity detector and parity generator circuits in the ACC;

Figures 15, 16, 17, and 18 show, respectively, the data flow during a Transaction Mode Normal Read, a Transaction Mode Failed Drive Read, a Transaction Mode Read-Modify-Write Read and a Transaction Mode Read-Modify-Write Write;

Fig. 19 is a schematic diagram of a set of disk drives in which check data is distributed among drives of the set according to a known technique;

Fig. 20 is a schematic diagram of a mass storage system suitable for use with the present invention;

Fig. 21 is a schematic diagram of the distribution of data on the surface of a magnetic disk;

Fig. 22 is a schematic diagram of the distribution of data in a first preferred embodiment of a redundancy group according to the present invention;

Fig. 23 is a schematic diagram of the distribution of data in a second, more particularly preferred embodiment of a redundancy group according to the present invention;

5 Fig. 24 is a diagram showing how the memory space of a device set might be configured in accordance with the principles of the present invention; and

Fig. 25 is a diagram of an exemplary embodiment of data structures for mapping between the logical levels of the present invention.

10

DETAILED DESCRIPTION OF THE DRAWINGS

The preferred embodiments of the present invention comprise a system for mass data storage. In the preferred embodiments described herein, the preferred devices for storing data are hard disk drives, referenced herein as disk drives. 15 Nothing herein should be understood to limit this invention to using disk drives only. Any other device for storing data may be used, including, but not limited to, floppy disks, magnetic tape drives, and optical disks.

20

Overall System Environment

One preferred embodiment of the present invention operates in the environment shown in Fig. 3. In Fig. 3, computer 10 communicates with a group of disk drive sets 18 through controller 11. In a preferred embodiment, controller 11 includes a number of components which permit computer 10 to access each of disk drive sets 18 even when there is a failure in one of the components of controller 11. As shown in Fig. 3, controller 11 includes a pair of two-level devices 13. Within 25 each of the two-level devices 13 is a first level controller 12 and a second level controller 14. A switch 16 which comprises a group of switches permits computer 10 to access disk drive sets 18 through more than one path. In this way, if either of 30 two-level devices 13 experience a failure in one of their components, the path may be re-routed without computer 10 being interrupted.

35 Fig. 3.5 is a diagram showing a pair of disk drive sets 18A and 18B connected to a pair of second level controllers 14A and 14B. Controllers 14A and 14B each include 40 two interface modules 27 for interfacing second level

controllers 14 with a pair of first level controllers 12 (shown in Fig. 3). Interface modules 27 are connected to buffers 330-335 which buffer data to be transmitted to and received from the disk drives. Second level controllers 14 are configured to 5 be primarily responsible for one group of disk drives and secondarily responsible for a second group of disk drives. As shown, second level controller 14A is primarily responsible for disk drives 20A1, 20B1, 20C1, 20D1 and 20E1 and secondarily responsible for disk drives 20A2, 20B2, 20C2, 20D2 and 20E2. A 10 spare drive 20X is shared by both second level controllers and is activated to take over for a disk drive which has failed.

The disk drives are connected to second level controllers 14 through a set of data interfaces 31. These interfaces are set by controllers 14 to configure the disk 15 drives in a particular arrangement. For example, disk drives 20A1, 20B1, 20C1 and 20D1 and 20A2, 20B2, 20C2 and 20D2 may be set to store data for the system while disk drives 20E1 and 20E2 may be set to store error correction codes. If any of the drives fail, drive 20X is set to take its place. Of course, 20 the disk drives of the system can be rearranged and may assume a wide variety of configurations.

Fig. 4 is more detailed diagram showing the interconnection of the components of the input/output system associated with computer 10 for accessing disk drive sets 18. 25 Computer 10 has its input/output ports connected to the first level controllers 12A and 12B. Second level controllers 14A and 14B are shown connected to first level controllers 12A and 12B. The lines between second level controllers 14 and first level controllers 12 represent data buses through which data 30 flows as well as control and status signals. The dashed line between second level controller 14A and second level controller 14B represents a communication line through which the second level controllers communicate with each other.

Second level controllers 14 are each connected to a 35 group of disk drive sets 18A-18F through switches 16A-16F.

Disk drives 20 are arranged in a manner so that each second level controller 14 is primarily responsible for one group of disk drive sets. As shown in Fig. 4, second level controller 14A may be primarily responsible for three of the 40 disk drive sets 18A-18F. Similarly, second level controller

14B may be primarily responsible for the remaining three disk drive sets 18A-18F. Second level controllers 14 are secondarily responsible for the disk drives primarily controlled by the partner second level controller. In the 5 particular arrangement shown in Fig. 4, second level controller 14A may be primarily responsible for the left three disk drive sets 18A, 18B and 18C and secondarily responsible for the right three disk drives sets 18D, 18E and 18F. Second level controller 14B is primarily responsible for the right three 10 disk drive sets 18D, 18E and 18F and secondarily responsible for the left three disk sets 18A, 18B and 18C.

Each second level controller 14 contains a second level controller recovery system (CRS) 22. CRS 22 is a portion of software code which manages the communication between second 15 level controllers 14 and first level controllers 12. CRS 22 is typically implemented as a state machine which is in the form of microcode or sequencing logic for moving second level controller 14 from state to state (described below). State changes are triggered as different events occur and messages 20 are sent between the various components of the system.

An ECC block 15 is also included in each second level controller 14. ECC block 15 contains circuitry for checking and correcting errors in data which occur as the data is passed between various components of the system. This circuitry is 25 described in more detail below.

Fig. 5 is a block diagram showing a more detailed illustration of the interconnections between second level controllers 14A and 14B and the disk drives. For simplicity, only a single disk drive port is shown. More disk drive ports 30 are included in the system as shown in Figs. 3 and 4.

Second level controller 14A has a primary control/sense line 50A for controlling its primary set of disk drives. An alternate control/sense line 52A controls an alternate set of disk drives. Of course, second level 35 controller 14B has a corresponding set of control/sense lines. Data buses 54A (second level controller 14A) and 54B (second level controller 14B) carry the data to and from disk drives 20. These data buses are typically in the form of a SCSI bus.

A set of switches 16A-16F are used to grant control of the disk drives to a particular second level controller. For example, in Fig. 4, second level controller 14A has primary responsibility for disk drives 20A-20C and alternate control of 5 disk drives 20D-20F. Second level controller 14B has primary control of disk drives 20D-20F and alternate control of disk drives 20A-20C. By changing the signals on control/sense lines 50 and 52, primary and secondary control can be altered.

Fig. 6 is a more detailed illustration of one of the 10 switches 16A-16F. A pair of pulse shapers 60A and 60B receive the signals from the corresponding control/sense lines 50A and 52B shown in Fig. 5. Pulse shapers 60 clean up the signals which may have lost clarity as they were transmitted over the lines. Pulse shapers of this type are well known in the art. 15 The clarified signals from pulse shapers 60 are then fed to the set and reset pins of R/S latch 62. The Q and \bar{Q} outputs of latch 62 are sent to the enable lines of a pair of driver/receivers 64A and 64B. Driver/receivers 64A and 64B are connected between the disk drives and second level controllers 20 14A and 14B. Depending upon whether primary control/sense line 52B or alternate control/sense line 50A is active, the appropriate second level controller will be in control at a particular time.

Fig. 7 is a state transition diagram showing the 25 relationships between the various states of CRS 22 (Fig. 3) of a particular second level controller 14. Each second level controller 14 must be in only one state at any particular point in time. Initially, assuming that the system is functioning properly and each second level controller 14 is primarily 30 responsible for half of the disk drive sets 18 and secondarily responsible for half of the disk drive sets 18, second level controller 14 is in a PRIMARY STATE 26. While in PRIMARY STATE 26, two major events may happen to move a second level controller 14 from PRIMARY STATE 26 to another state. The 35 first event, is the failure of the particular second level controller 14. If there is a failure, second level controller 14 shifts from PRIMARY STATE 26 to a NONE STATE 28. In the

process of doing so, it will pass through RUN-DOWN-PRIMARIES-TO-NONE STATE 30.

There are two types of failures which are possible in second level controller 14. The first type of failure is a controlled failure. Further, there are two types of controlled failures.

The first type of controlled failure is a directed controlled failure. This is not actually a failure but instead an instruction input from an outside source instructing a particular second level controller to shut down. This instruction may be received in second level controller 14 from one of the following sources: An operator, through computer 10; a console 19 through a port 24 (e.g. RS-232) on the first level controller; a diagnostic console 21 through a port 23 (e.g. RS-232) on the second level controller; or by software initiated during predictive maintenance. Typically, such an instruction is issued in the case where diagnostic testing of a second level controller is to be conducted. In a directed controlled failure, the second level controller finishes up any instructions it is currently involved with and refuses to accept any further instructions. The second level controller effects a "graceful" shut down by sending out messages to the partner second level controller that it will be shutting down.

The second type of controlled failure is referred to as a moderate failure. In this case, the second level controller recognizes that it has a problem and can no longer function properly to provide services to the system. For example, the memory or drives associated with that second level controller may have malfunctioned. Therefore, even if the second level controller is properly functioning, it cannot adequately provide services to the system. It aborts any current instructions, refuses to accept any new instructions and sends a message to the partner second level controller that it is shutting down. In both controlled failures, the malfunctioning second level controller releases the set of disk drives over which it has control. These drives are then taken over by the partner second level controller.

The second type of failure is a complete failure. In a complete failure, the second level controller becomes inoperable and cannot send messages or "clean-up" its currently pending instructions by aborting them. In other words, the 5 second level controller has lost its ability to serve the system. It is up to one of the first level controllers or the partner second level controller to recognize the problem. The partner second level controller then takes control of the drives controlled by the malfunctioning second level 10 controller. The routing through the malfunctioning second level controller is switched over to the partner second level controller.

In all of the above failures, the switching takes place without interruption to the operation of the computer. 15 Second level controllers 14 and first level controllers 12 handle the rerouting independently by communicating the failure among themselves.

Assuming there was a failure in second level controller 14A, second level controller 14A moves from PRIMARY 20 STATE 26 through a transition RUN-DOWN-PRIMARIES-TO-NONE STATE 30 to NONE STATE 28. At the same time, properly functioning second level controller 14B moves from PRIMARY STATE 26 to BOTH STATE 32. The basis for the change in state of each of second level controllers 14A and 14B is the failure of second level 25 controller 14A. When a second level controller fails, it is important to switch disk drive control away from the failed second level controller. This permits computer 10 to continue to access disk drives which were formerly controlled by a particular second level controller which has failed. In the 30 current example (Fig. 4), disk drive sets 18A-18C are switched by switching functions 16A-16C so that they are controlled by second level controller 14B. Therefore, second level controller 14B is in BOTH STATE 32 indicating that it has control of the disk drive sets 18 for both second level 35 controllers. Second level controller 14A now controls none of the disk drives and is in NONE STATE 28. The transition state 30 determines which of several possible transition paths is used.

If second level controller 14A is in NONE STATE 28 and second level controller 14B is in BOTH STATE 32 there are a number of options for transferring control of disk drive sets 18A-18F once second level controller 14A has been repaired.

- 5 First, second level controller 14A and second level controller 14B could each be shifted back to PRIMARY STATE 26. This is accomplished for drive sets 18A-18C by second level controller 14A moving from NONE STATE 28 directly to PRIMARY STATE 26 along the preempt p line. Preempt p simply stands for "preempt primary" which means that second level controller 14A preempts its primary drives or takes control of them from second level controller 14B. At the same time, second level controller 14B moves from BOTH STATE 32 through a transition RUN-DOWN-SECONDARIES-TO-PRIMARIES STATE 34 and then to PRIMARY STATE 26.

- 15 A second alternative is for second level controller 14A to move from NONE STATE 28 to SECONDARY STATE 36. Once in SECONDARY STATE 36, second level controller 14A is in control of its secondary disk drive sets 18D-18F. Second level controller 14B concurrently moves from BOTH STATE 32 through RUN-DOWN-PRIMARIES-TO-SECONDARIES STATE 38 and on to SECONDARY STATE 36. When both second level controllers are in SECONDARY STATE 36, they are in control of their secondary disk drive sets. Second level controller 14A controls disk drive sets 18D-18F and second level controller 14B controls disk drive sets 18A-18C.

- 20 From SECONDARY STATE 36, a failing second level controller 14 may move through RUN-DOWN-SECONDARIES-TO-NONE STATE 40 to NONE STATE 28. If this occurs, the properly functioning partner second level controller 14 moves from SECONDARY STATE 36 to BOTH STATE 32 so that computer 10 could access any one of disk drive sets 18. As in the previous example, if second level controller 14A were to fail it moves from SECONDARY STATE 36 through RUN-DOWN-SECONDARIES-TO-NONE STATE 40 and into NONE STATE 28. At the same time, properly functioning second level controller 14B moves from SECONDARY STATE 36 along the preempt b/p line into BOTH STATE 32. Preempt b/p stands for "preempt both/primaries." In other

words, all of the disk drives are preempted by the properly functioning second level controller.

If, for all sets 18, second level controller 14A is in NONE STATE 28 and second level controller 14B is in BOTH STATE 32, it is possible for second level controller 14A to take control of all sets 18 of disk drives. This is desirable if second level controller 14A were repaired and second level controller 14B failed. Second level controller 14A moves from NONE STATE 28 along the preempt b line to BOTH STATE 32. At the same time, second level controller 14B moves from BOTH STATE 32 through RUN-DOWN-BOTH-TO-NONE STATE 42 and into NONE STATE 28. At this point, second level controller 14A controls all disk drives while second level controller 14B controls none of the disk drives.

Various failures may trigger the movement of second level controllers 14 between states. Between states a number of events take place. Each of these events is described in Figs. 8A-8I. In Fig. 8A, second level controller 14 is in PRIMARY STATE 26. There are three different events which can take place while second level controller 14 is in PRIMARY STATE 26. The first event is for a preempt message 100 to be received from the partner second level controller. At this point, the second level controller receiving such a message will take the secondary path, represented by block 102, and end up at BOTH STATE 32. The second path which may be taken is triggered by receipt of a message 104 from CRS 22 of the other second level controller. This may be some sort of communication which results in the second level controller remaining in PRIMARY STATE 26. It will report and return messages 106 to the other second level controller. The final path which may be taken results in second level controller ending up in RUN-DOWN-PRIMARIES-TO-NONE STATE 30. This path is triggered upon receipt of a message 108 to release both sets of drives or the primary disk drives. A timer is then set in block 110 and upon time out a message 112 is sent to the other second level controller to take control of the primary set of disk drives. Once in RUN-DOWN-PRIMARIES-TO-NONE STATE 30,

second level controller 14 will eventually end up in NONE STATE 28.

Fig. 8B illustrates various paths from RUN-DOWN-PRIMARIES-TO-NONE STATE 30 to NONE STATE 28. Three possible events may take place. First, a message 114 may be received from another second level controller providing communication information. In this case, second level controller 14 reports back messages 116 and remains in RUN-DOWN-PRIMARIES-TO-NONE STATE 30. The second event which may occur is for the timer, set during transition from PRIMARY STATE 26 to RUN-DOWN-PRIMARIES-TO-NONE STATE 30 to time out 118. If this happens, second level controller 14 realizes that message 112 (Fig. 8A) didn't get properly sent and that there has been a complete failure. It releases control of both its primaries and secondary disk drives 122. It then ends up in NONE STATE 28. The third event which may occur while in RUN-DOWN-PRIMARIES-TO-NONE STATE 30 is for a response to be received 124 from message 112 (Fig. 8A) sent out while second level controller moved from PRIMARY STATE 26 to RUN-DOWN-PRIMARIES-TO-NONE STATE 30. This response indicates that the message was properly received. Second level controller 14 then releases its primary drives 126 and ends up in NONE STATE 28.

Fig. 8C covers the state transition between NONE STATE 28 and one of either BOTH STATE 32, PRIMARY STATE 26, or SECONDARY STATE 36. When in NONE STATE 28, second level controller 14 can only receive messages. First, it may receive a message 128 instructing it to preempt both its primary and alternative sets of disk drives. It performs this function 130 and ends up in BOTH STATE 32. A second possibility is for it to receive a preempt message 132 instructing it to preempt its primary set of drives. It performs this instruction and ends up in PRIMARY STATE 26. A third alternative is the receipt of a preempt message 136 instructing second level controller 14 to preempt its secondary drives. Upon performance of this instruction 138 it ends up in SECONDARY STATE 36. Finally, while in NONE STATE 28 second level controller 14 may receive communication messages 140 from its partner second level

controller. It reports back 142 to the other second level controller and remains in NONE STATE 28.

Fig. 8D illustrates the movement of second level controller 14 from SECONDARY STATE 36 to BOTH STATE 32 or RUN-DOWN-SECONDARIES-TO-NONE STATE 40. While in SECONDARY STATE 36, any one of three messages may be received by second level controller 14. A first possibility is for a preempt both or primary message 144 to be received. At this point, second level controller 14 takes control of its primary drives 146 and ends up in BOTH STATE 32. A second possibility is for communication messages 148 to be received from the partner controller. This results in second level controller 14 reporting back 150 and remaining in its present SECONDARY STATE 36. Finally, a release both or secondary message 152 may be received. Second level controller 14 sets a timer 154 upon receipt of this message. It then sends out a message 156 indicating it is now in RUN-DOWN-SECONDARIES-TO-NONE STATE 40.

Fig. 8E shows the transition of second level controller 14 from RUN-DOWN-SECONDARIES-TO-NONE STATE 40 to NONE STATE 28. Three different messages may be received during RUN-DOWN-SECONDARIES-TO-NONE STATE 40. First, messages 158 from the partner second level controller may be received. Second level controller 14 then reports back (160) to its partner and remains in RUN-DOWN-SECONDARIES-TO-NONE STATE 40. A second possibility is for the timer, set between SECONDARY STATE 36 and the present state, to time out (162). This indicates that message 156 (Fig. 8D) was not properly sent out and received by the partner second level controller and that there has been a complete failure to second level controller 14. Second level controller 14 then reports out (164) that it will release both of its sets of disk drives 166. This results in it moving to NONE STATE 28. Finally, second level controller 14 may receive a response 168 to its message 156 (Fig. 8D) sent after setting the timer between SECONDARY STATE 36 and RUN-DOWN-SECONDARIES-TO-NONE STATE 40. Upon receiving this response, it releases its secondary drives and ends up in NONE STATE 28.

Fig. 8F illustrates the various paths from BOTH STATE 32 to any one of RUN-DOWN-PRIMARIES-TO-SECONDARIES STATE 38, RUN-DOWN-SECONDARIES-TO-PRIMARIES STATE 34 or RUN-DOWN-BOTH-TO-NONE STATE 42. A first possible message which may be received during BOTH STATE 32 is a release primary message 172. This will cause second level controller 14 to set a timer 174, send a message 176 indicating it is running down primaries, and wait in RUN-DOWN-PRIMARIES-TO-SECONDARIES STATE 38. A second message which may be received is a release secondaries message 180. Upon receiving release secondaries message 180, second level controller 14 sets a timer 182 and sends a message 184 indicating it has moved into RUN-DOWN-SECONDARIES-TO-PRIMARIES STATE 34. A third possibility for second level controller 14 is to receive communication messages 186 from its partner second level controller. It will report back (188) and remain in BOTH STATE 32. Finally, second level controller 14 may receive an instruction 190 telling it to release both primary and secondary sets of drives. At this point it sets the timer 192 and sends out a message 194 that it has released both primary and secondary drive sets. It will then remain in the RUN-DOWN-BOTH-TO-NONE STATE 42 until it receives further instructions from the other second level controller.

Fig. 8G shows the various paths by which second level controller 14 moves from RUN-DOWN-PRIMARIES-TO-SECONDARIES STATE 38 to one of either NONE STATE 28 or SECONDARY STATE 36. The first possibility is that second level controller 14 receives messages 196 from the other second level controller. It then reports back (198) and remains in RUN-DOWN-PRIMARIES-TO-SECONDARIES STATE 38. A second possibility is that the timer (174), set between BOTH STATE 32 and RUN-DOWN-PRIMARIES-TO-SECONDARIES STATE 38 times out (200). At this point, second level controller 14 realizes that message 176 (Fig. 8F) was not properly sent. A complete failure has occurred. The second level controller reports (202) that it has released both sets of disk drives, and releases both sets (204). Second level controller 14 then enters NONE STATE 28. Finally, a run down path response message 206 is received acknowledging receipt of message 176 (Fig. 8F) sent between BOTH STATE 32 and RUN-DOWN-

PRIMARIES-TO-SECONDARIES STATE 38. Second level controller 14 releases its primary drives 208 and enters SECONDARY STATE 36.

Fig. 8H shows the possible paths down which second level controller 14 moves between RUN-DOWN-SECONDARIES-TO-

- 5 PRIMARIES STATE 34 and one of either NONE STATE 28 or PRIMARY STATE 26. A first possibility is that second level controller 14 receives a message 210 from the other second level controller. It then reports back (212) and remains in RUN-DOWN-SECONDARIES-TO-PRIMARIES STATE 34. A second possibility
- 10 is that the timer (182), set between BOTH STATE 32 and RUN-DOWN-SECONDARIES-TO PRIMARY-STATE 34 times out (214). If this occurs, second level controller 14 realizes that message 184 (Fig. 8F) was not properly sent. A complete failure has occurred. Second level controller then sends a message 216
- 15 indicating that it has released its drives and then it releases both primary and secondary disk drive sets (218) which it controls. Second level controller then moves into NONE STATE 28. Finally, a third possibility is that second level controller 14 receives a response 220 to message 184 (Fig. 8F)
- 20 sent between BOTH STATE 32 and RUN-DOWN-SECONDARIES-TO-PRIMARIES-STATE 34. It will then release (222) its secondary drives and enter PRIMARY STATE 26.

Fig. 8I shows the possible paths illustrating the transition of second level controller between RUN-DOWN-BOTH-

- 25 TO-NONE STATE 42 and NONE STATE 28. Three possible events may take place. First, a message 230 may be received from the other second level controller providing communication information. In this case, second level controller 14 reports back messages 232 and remains in RUN-DOWN-BOTH-TO-NONE STATE 42. The second event which may occur is for the timer (192), set during transition from BOTH STATE 32 to RUN-DOWN-BOTH-TO-NONE STATE 42, to time out (234). If this happens, second level controller 14 realizes that message 194 (Fig. 8F) sent during BOTH STATE 32 didn't get properly sent and that there
- 30 has been a complete failure. It releases control of both its primaries and secondary disk drives (238). It then ends up in NONE STATE 28. The third event which may occur while in RUN-DOWN-BOTH-TO-NONE STATE 42 is for a response to be received
- 35

(240) from message 194 (Fig. 8F) sent out while second level controller moved from BOTH STATE 32 to RUN-DOWN-BOTH-TO-NONE STATE 42. This response indicates that the message was properly received. Second level controller 14 then releases 5 both sets of drives (242) and ends up in NONE STATE 28.

Rerouting Data Paths Between Buffers and Disk Drives

Figure 9 illustrates a first preferred embodiment of circuitry for rerouting data paths between buffers and disk drives 20. In Figure 9, X-bar switches 310 through 315 are coupled to a bus 309 communicating with the second level controller engine (see Figures 3 and 4). In turn, each X-bar switch is coupled by a bus to disk drives 20A1 through 20A6 and to each buffer 330 through 336. Bus 350 couples each buffer to 10 a first level controller which are coupled to a computer such as computer 10 (Figures 3 and 4). In this embodiment, although only six disk drives are illustrated, any arbitrary number could be used, as long as the illustrated architecture is preserved by increasing the number of X-bar switches and output 15 buffers in a like manner and maintaining the interconnected bus structures illustrated in Figure 9.

In operation, the second level controller will load various registers (not illustrated herein) which configure the X-bar switches to communicate with particular buffers and 20 particular disk drives. The particular configuration can be changed at any time while the system is operating. Data flow is bi-directional over all the buses. By configuring the X-bar switches, data flowing from any given buffer may be sent to any given disk drive or vice versa. Failure of any particular 25 system element does not result in any significant performance degradation, as data flow can be routed around the failed element by reconfiguring the registers for the X-bar switch. In a preferred mode of operation, data may be transferred from or to a particular disk drive in parallel with other data transfers occurring in parallel on every other disk drive. 30 This mode of operation allows for a very high input/output rate of data as well as a high data throughput.

To illustrate this embodiment's mode of operation, the following example is offered. Referring to Figure 9, assume that all data flow is initially direct, meaning, for example, that data in buffer 330 flows directly through X-bar switch 110 to disk drive 20A1. Were buffer 330 to fail, the registers of X-bar switch 310 could be reconfigured, enabling X-bar switch 310 to read data from buffer 335 and direct that data to disk drive 20A1. Similar failures in other buffers and in the disk drives could be compensated for in the same manner.

10

Generation of Redundancy Terms and Error Detection on Parallel Data

Figure 10 illustrates a second preferred embodiment of the present invention. This second embodiment incorporates Array Correction Circuits ("ACCs") to provide error detection and correction capabilities within the same general architecture as illustrated for the first preferred embodiment shown in Figure 9. To ease the understanding of this embodiment, the full details of the internal structure of both the X-bar switches (310 through 315) and the ACC circuits 360 and 370 are not shown in Fig 10. Figures 11 and 12 illustrate the internal structure of these devices and will be referenced and discussed in turn. Additionally, bus LBE as illustrated in Figure 10 does not actually couple the second level controller (Figures 3 and 4) directly to the X-bar switches, the ACCs, and the DSI units. Instead, the second level controller communicates with various sets of registers assigned to the X-bar switches, the ACCs and the DSI units. These registers are loaded by the second level controller with the configuration data which establishes the operating modes of the aforementioned components. As such registers are known, and their operation incidental to the present invention, they are not illustrated or discussed further herein.

35

The embodiment shown in Figure 10 shows data disk drives 20A1 through 20A4 and P and Q redundancy term drives 20A5 and 20A6. A preferred embodiment of the present invention utilizes 13 disk drives: ten for data, two for P and Q redundancy terms, and one spare or backup drive. It will be

understood that the exact number of drives, and their exact utilization may vary without in any way changing the present invention. Each disk drive is coupled by a bi-directional bus (Small Computer Standard Interface) to units 340 through 345, 5 herein labelled DSI. The DSI units perform some error detecting functions as well as buffering data flow into and out of the disk drives.

Each DSI unit is in turn coupled by a bi-directional bus means to an X-bar switch, the X-bar switches herein numbered 310 through 315. The X-bar switches are coupled in turn to word assemblers 350 through 355 by means of a bi-directional bus. The bus width in this embodiment is 9 bits, 8 for data, 1 for a parity bit. The word assemblers assemble 36-bit (32 data and 4 parity) words for transmission 10 to buffers 330 through 335 over bi-directional buses having a 36-bit width. When data flows from the output buffers to the X-bar switches, the word assemblers decompose the 36-bit words 15 into the 9-bits of data and parity.

The X-bar switches are also coupled to ACC units 348 and 349. The interconnection between the X-bar switches and the ACCs is shown in more detail in Fig. 11. Each X-bar switch can send to both or either ACC the 8 bits of data and 1 parity bit that the X-bar switch receives from either the DSI units or the word assemblers. In turn, the X-bar switches can receive 20 9 bits of the P and Q redundancy terms calculated by the ACCs over lines E_1 and E_2 . As shown, the ACCs can direct the P and Q redundancy terms to any X-bar switch, not being limited to the disk drives labelled P and Q. Depending on the configuration commanded by the second level controller, ACCs 348 and 349 can 25 be mutually redundant, in which case the failure of one or the other ACC does not affect the system's ability to detect or correct errors, or each ACC can detect and correct errors on a portion of the total set of disk drives. When operating in this second manner, certain specific types of operations which 30 write data to individual disk drives are expedited, as each ACC can write to a separate individual disk drive. The specific disk drives that the individual ACCs monitor can be 35 reconfigured at any time by the second level controller.

The illustrated connections of the ACCs and the X-bar switches also allows data to be switched from any X-bar switch to any ACC once the second level controller configures the related registers. This flexibility allows data to be routed
5 away from any failed disk drive or buffer.

Figure 11 shows important internal details of the ACCs and the X-bar switches. X-bar switch 310 is composed of two mirror-image sections. These sections comprise, respectively, 9-bit tri-state registers 370/380,
10 multiplexers 372/382, first 9-bit registers 374/384, second 9-bit registers 376/386, and input/output interfaces 379/389. In operation, data can flow either from the word assembler to the DSI unit or vice versa.

Although many pathways through the X-bar switch are possible, as shown by Figure 11, two aspects of these pathways are of particular importance. First, in order to allow the ACC sufficient time to calculate P and Q redundancy terms or to detect and correct errors, a data pathway of several registers can be used, the data requiring one clock cycle to move from
15 one register to the next. By clocking the data through several registers, a delay of sufficient length can be achieved. For example, assuming a data flow from the word assembler unit to a disk drive, 9 bits are clocked into 9-bit register 374 and tri-state register 370 on the first clock pulse. On the next clock
20 pulse, the data moves to 9-bit register 386 and through redundancy circuit 302 in the ACC 348 to P/Q registers 304 and 306. The next clock pulses move the data to the DSI unit.
25

The second important aspect of the internal pathways relates to the two tri-state registers. The tri-state
30 registers are not allowed to be active simultaneously. In other words, if either tri-state register 370 or 380 is enabled, its counterpart is disabled. This controls data transmission from the X-switch to the ACC. The data may flow only from the DSI unit to the ACC or from the word assembler to the ACC, but not from both to the ACC simultaneously. In the opposite direction, data may flow from the ACC to the word assembler and the DSI simultaneously.
35

ACC unit 348 comprises a redundancy circuit 302, wherein P and Q redundancy terms are generated, P and Q registers 304 and 306, wherein the P and Q redundancy terms are stored temporarily, regenerator and corrector circuit 308, 5 wherein the data from or to a failed disk drive or buffer can be regenerated or corrected, and output to interfaces 390, 391, 392 and 393.

Redundancy Generation and Error Checking Equations

10 The main functional components of the second preferred embodiment and their physical connections to one another have now been described. The various preferred modes of operation will now be described. In order to understand these functional modes, some understanding of the error 15 detection and correction method used by the present invention will be necessary.

Various error detection and correction codes are known and used in the computer industry. Error-Control Coding and Applications, D. Wiggert, The MITRE Corp., describes 20 various such codes and their calculation. The present invention in this second preferred embodiment is implemented using a Reed-Solomon error detection and correction code. Nothing herein should be taken to limit the present invention to using only a Reed-Solomon code. If other codes were used, 25 various modifications to the ACCs would be necessary, but these modifications would in no way change the essential features of this invention.

Reed-Solomon codes are generated by means of a field generator polynomial, the one used in this embodiment being 30 $x^4 + x + 1$. The code generator polynomial needed for this Reed-Solomon code is $(x + a^0) \cdot (x + a^1) = x^2 + a^4x + a^1$. The generation and use of these codes to detect and correct errors is known.

The actual implementation of the Reed-Solomon code in 35 the present invention requires the generation of various terms and syndromes. For purposes of clarity, these terms are generally referred to herein as the P and Q redundancy terms. The equations which generate the P and Q redundancy terms are:

$$P = d_{n-1} + d_{n-2} + \dots + d_1 + d_0$$

5

and

$$Q = d_{n-1} \cdot a_{n-1} + d_{n-2} \cdot a_{n-2} + \dots + d_1 \cdot a_1 + d_0 \cdot a_0.$$

- 10 The P redundancy term is essentially the simple parity of all the data bytes enabled in the given calculation. The Q logic calculates the Q redundancy for all data bytes that are enabled. For Q redundancy, input data must first be multiplied by a constant "a" before it is summed. The logic operations
 15 necessary to produce the P and Q redundancy terms are shown in Figures 12a and 12b. All operations denoted by \oplus are exclusive-OR ("XOR") operations. Essentially, the final P term is the sum of all P_i terms. The Q term is derived by multiplying all Q_i terms by a constant and then XORing the results.
 20 These calculations occur in redundancy circuit 302 in ACC 260 (Figure 11). The second preferred embodiment, using its implementation of the Reed-Solomon code, is able to correct the data on up to two failed disk drives.

The correction of data requires the generation of
 25 additional terms S_0 and S_1 within the ACC. Assuming that the P and Q redundancy terms have already been calculated for a group of data bytes, the syndrome equations

30

$$S_0 = d_{n-1} + d_{n-2} + \dots + d_1 + d_0 + P$$

35

$$S_1 = (d_{n-1} \cdot a_{n-1}) + (d_{n-2} \cdot a_{n-2}) + \dots + (d_1 \cdot a_1) + (d_0 \cdot a_0) + Q$$

are used to calculate S_0 and S_1 . For S_0 an ACC register enables the necessary data bytes and the P redundancy to be used in the calculation. For S_1 , the necessary input data must 40 first be multiplied by a_i before being summed with the Q redundancy information.

As stated, an ACC can correct the data on up to two failed disk drives in this embodiment. The failed disk drive register (not illustrated) in the relevant ACC will be loaded

30

with the address of the failed disk or disks by the second level controller. A constant circuit within the ACC will use the drive location information to calculate two constants k_0 and k_1 as indicated in Table 1 below, where i represents the 5 address of the first failed disk drive, j is the address of the second failed disk drive, and a is a constant. The columns labelled Failed Drives indicate which drives have failed. Column k_0 and k_1 indicate how those constants are calculated given the failure of the drives noted in the Failed Drives 10 columns.

TABLE 1

15

	<u>Failed Drives</u>	k_0	k_1
20	P	-	0
	Q	-	1
	i	-	0
	i	P	0
	Q	i	0
25	i	j	$aj/ai+aj$
	P	Q	0

The error correction circuits use the syndrome information S_0 30 and S_1 , as well as the two constants k_0 and k_1 to generate the data contained on the failed disk drives. The error correction equations are as follows:

$$F_1 = S_0 \cdot k_0 + S_1 \cdot k_1$$

35

$$F_2 = S_0 + E_1$$

F_1 is the replacement data for the first failed disk drive. F_2 is the replacement data for the second failed disk drive. The 40 equations which generate the P and Q redundancy terms are realized in combinatorial logic, as is partially shown in

Figures 12a and 12b. This has the advantage of allowing the redundancy terms to be generated and written to the disk drives at the same time that the data is written to the drives. This mode of operation will be discussed later.

5

Operational Modes

Having described the aspects of the Reed-Solomon code implementation necessary to understand the present invention, the operational modes of the present invention will now be discussed.

10

The second preferred embodiment of the present invention operates primarily in one of two classes of operations. These are parallel data storage operations and transaction processing operation. These two classes of 15 operations will now be discussed with reference to the figures, particularly Figures 10, 13 and 14 and Tables 2 through 7.

Although Figure 10 only shows four data drives and the P and Q redundancy term drives, a preferred embodiment uses a set of 13 disk drives, 10 for data, 2 for the P and Q terms, 20 and a spare. Although nothing herein should be construed to limit this discussion to that specific embodiment, parallel processing operations will be described with relation to that environment.

25

Parallel Processing Operations

In parallel processing operations, all the drives are considered to comprise a single large set. Each of the disk drives will either receive or transmit 9 bits of data simultaneously. The result of this is that the 9-bits of data 30 appearing in the DSI units of all the drives simultaneously are treated as one large codeword. This result is shown in Figure 13a. Codeword 400 comprises 9 bits of data from or for disk drive d_{n-1} , 9 bits of data from or for disk drive d_{n-2} , and so on, with the P and Q disk drives receiving or 35 transmitting the P and Q redundancy term. In a parallel write operation, all the disk drives in the set, except for the spare disk drive, will receive a byte of data (or a redundancy term whose length is equal to the data byte) simultaneously. As

shown, the same sector in all the disk drives will receive a part of codeword 400. For example, in the illustration, sector 1 of disk drive n-1 will receive a byte of data designated d_{n-1} from codeword 400, sector 1 of disk drive n-2 will receive a 5 byte of data designated d_{n-2} from codeword 400 and so on.

In the actual implementation of this preferred embodiment, the codewords are "striped" across the various disk drives. This means that for each successive codeword, different disk drives receive the P and Q redundancy terms. In 10 other words, drive d_{n-1} is treated as drive d_{n-2} for the second codeword and so on, until what was originally drive d_{n-1} receives a Q redundancy term. Thus, the redundancy terms "stripe" through the disk drives.

15 Pairs of P and Q Terms for Nibbles

Calculating the P and Q redundancy terms using 8-bit symbols would require a great deal of hardware. To reduce this hardware overhead, the calculations are performed using 4-bit bytes or nibbles. This hardware implementation does not change 20 the invention conceptually, but does result in the disk drives receiving two 4-bit data nibbles combined to make one 8-bit byte. In Figure 13b, codeword 450, as well as the illustrated sectors A of the disk drives, illustrate how the codeword is broken up and how the disk drives receive upper and lower 4-bit nibbles. Table 2 shows how, for codewords one through N, a 25 different portion of the codeword is placed on the different drives. Each disk drive, for a given codeword, receives an upper and lower 4-bit nibble, designated with L's and U's, of the codeword. Additionally, the same section is used to store 30 the nibbles on each of the disk drives used to store the codeword. In other words, for codeword₁, the first sector of disk drives n-1 through 0 receives the nibbles.

TABLE 2

5

CODEWORD -- DATA AND P AND Q

10

Sector of
Drive d_{n-1}Sector of
Drive d_{n-2}Sector of
Drive d₀Sector of
Drive PSector of
Drive Q

15

Codeword₁Codeword₁Codeword₁Codeword₁Codeword₁Codeword₁(d_{n-1_L})(d_{n-1_U})(d_{n-2_L})(d_{n-2_U})(d_{0_L})(d_{0_U})(P_{1_L})(P_{1_U})(Q_{1_L})(Q_{1_U})

20

Codeword₂Codeword₂Codeword₂Codeword₂Codeword₂Codeword₂(d_{n-1_L})(d_{n-1_U})(d_{n-2_L})(d_{n-2_U})(d_{0_L})(d_{0_U})(P_{2_L})(P_{2_U})(Q_{2_L})(Q_{2_U})

25

Codeword_nCodeword_nCodeword_nCodeword_nCodeword_nCodeword_n(d_{n-1_L})(d_{n-1_U})(d_{n-2_L})(d_{n-2_U})(d_{0_L})(d_{0_U})(P_{n_L})(P_{n_U})(Q_{n_L})(Q_{n_U})

30

Referring back to Figure 10, for a parallel data write to the disks, the data is provided in parallel from buffers 330, 331, 332 and 333 along those data buses coupling the buffers to X-bar switches 310, 311, 312, and 313 after the 36-bits of data are disassembled in word assemblers 350 through 353 into 9-bit words. These X-bar switches are also coupled to inputs D3, D2, D1 and D0, respectively, of ACC 348 and ACC 349.

In parallel processing modes, the two ACCs act as mutual "backups" to one another. Should one fail, the other will still perform the necessary error correcting functions. In addition to operating in a purely "backup" condition, the second level controller engine configures the ACCs so that each ACC is performing the error detection and correction functions for a portion of the set, the other ACC performing these functions for the remaining disk drives in the set. As the ACC units are still coupled to all the disk drives, failure of one

or the other unit does not impact the system as the operating ACC can be reconfigured to act as the dedicated ACC unit for the entire set. For purposes of discussion, it is assumed here that ACC 348 is operating. ACC 348 will calculate the P and Q

- 5 redundancy term for the data in the X-bar switches and provide the terms to its E_1 and E_2 outputs, which outputs are coupled to all the X-bar switches. For discussion only, it is assumed that only the E_2 connection of X-bar switch 314 and the E_1 connection of X-bar switch 315 are enabled. Thus, although the
- 10 data is provided along the buses coupling ACC 348's E_1 and E_2 output to all the X-bar switches, the Q term is received only by X-bar switch 314 and the P term is received by X-bar switch 315. Then, the Q and P terms are provided first to DSI units 344 and 345 and then disk drives 20A5 and 20A6. It
- 15 should be recalled that the various internal registers in the X-bar switches will act as a multi-stage pipeline, effectively slowing the transit of data through the switches sufficiently to allow ACC 348's redundancy circuit 302 to calculate the P and Q redundancy terms.

- 20 As ACC 349 is coupled to the X-bar switches in a substantially identical manner to ACC 348, the operation of the system when ACC 349 is operational is essentially identical to that described for ACC 348.

- Subsequent parallel reads from the disks occur in the
- 25 following manner. Data is provided on bi-directional buses to DSI units 340, 341, 342 and 343. P and Q redundancy terms are provided by DSI units 345 and 344, respectively. As the data and P and Q terms are being transferred through X-bar switches 310 through 315, ACC 348 uses the P and Q terms to determine if
- 30 the data being received from the disk drives is correct. Word assemblers 350 through 353 assemble successive 9-bit words until the next 36-bits are available. This 36-bits are forwarded to buffers 330 through 333. Note that the 9-bit words are transmitted to the buffers in parallel. If that data
- 35 is incorrect, the second level controller will be informed.

During a parallel read operation, in the event that there is a failure of a disk drive, the failed disk drive will, in certain instances, communicate to the second level

controller that it has failed. The disk drive will communicate with the second level controller if the disk drive cannot correct the error using its own corrector. The second level controller will then communicate with ACCs 348 and 349 by 5 loading the failed drive registers in the ACC (not shown in the figures) with the address of the failed drive. The failed drive can be removed from the set by deleting its address from the configuration registers. One of the set's spare drives can then be used in place of the failed drive by inserting the 10 address of the spare drive into the configuration registers.

The ACC will then calculate the replacement data necessary to rewrite all the information that was on the failed disk onto the newly activated spare. In this invention, the term spare or backup drive indicates a disk drive which 15 ordinarily does not receive or transmit data until another disk drive in the system has failed.

When the data, P, and Q bytes are received, the ACC circuits use the failed drive location in the failed drive registers to direct the calculation of the replacement data for 20 the failed drive. After the calculation is complete, the data bytes, including the recovered data, are sent to data buffers in parallel. Up to two failed drives can be tolerated with the Reed-Solomon code implemented herein. All operations to replace failed disk drives and the data thereon occur when the 25 system is operating in a parallel mode.

Regeneration of data occurs under second level controller control. When a failed disk drive is to be replaced, the ACC regenerates all the data for the replacement disk. Read/write operations are required until all the data 30 has been replaced. The regeneration of the disk takes a substantial amount of time, as the process occurs in the background of the system's operations so as to reduce the impact to normal data transfer functions. Table 3 below shows the actions taken for regeneration reads. In Table 3, i 35 represents a first failed drive and j represents a second failed drive. In Table 3, the column labelled Failed Drives indicates the particular drives that have failed. The last

column describes the task of the ACC given the particular indicated failure.

5

TABLE 3

Regeneration Read

10

Failed Drives

	P	-	ACC calculates P redundancy
	Q	-	ACC calculates Q redundancy
15	i	-	ACC calculates replacement data for i drive
	i	P	ACC calculates replacement data for i drive and P redundancy
20	Q	i	ACC calculates replacement data for i drive and Q redundancy
	j	i	ACC calculates replacement data for i and j drives
	P	Q	ACC calculates P and Q redundancy

25

It should be noted that if both a data disk drive and a redundancy disk drive fail, the data on the data disk drive must be regenerated before the redundancy terms on the redundancy drive. During a regeneration write, regeneration data or redundancy terms are written to a disk and no action is required from the ACC logic.

During a parallel read operation, it should also be noted that additional error detection may be provided by the ACC circuitry.

35

Table 4 indicates what actions may be taken by the ACC logic unit when the indicated drive(s) has or have failed during a failed drive read operation. In this operation, the drives indicated in the Failed Drives columns are known to have failed prior to the read operation. The last column indicates the ACC response to the given failure.

TABLE 4

5

Failed Drives

10

P		No action by ACC
Q		No action by ACC
i		ACC calculates replacement data
i	P	ACC calculates the replacement data
Q	i	ACC calculates the replacement data
i	j	ACC calculates replacement data
P	Q	No action by ACC

15

Transaction Processing Mode: Read

Transaction processing applications require the ability to access each disk drive independently. Although each disk drive is independent, the ACC codeword with P and Q redundancy is maintained across the set in the previously described manner. For a normal read operation, the ACC circuitry is not generally needed. If only a single drive is read, the ACC cannot do its calculations since it needs the data from the other drives to assemble the entire codeword to recalculate P and Q and compare it to the stored P and Q. Thus, the data is assumed to be valid and is read without using the ACC circuitry (see Figure 15). Where drive 20C1 is the one selected, the data is simply passed through DSI unit 342 X-bar switch 312, word assembler 352 and buffer 332 to the external computer. If the disk drive has failed, the read operation is the same as a failed drive read in parallel mode with the exception that only the replacement data generated by the ACC is sent to the data buffer. In this case, the disk drive must notify the second level controller that it has failed, or the second level controller must otherwise detect the failure. Otherwise, the second level controller will not know that it should read all the drives, unless it assumes that there might be an error in the data read from the desired drive. The failed drive read is illustrated in Figure 16, with drive 20C1

having the desired data, as in the example of Figure 15. In Figure 16, the second level controller knows that drive 20C1 has failed, so the second level controller calls for a read of all drives except the failed drive, with the drive 20C1 data

- 5 being reconstructed from the data on the other drives and the P and Q terms. Only the reconstructed data is provided to its buffer, buffer 332, since this is the only data the external computer needs.

10 Transaction Processing Mode: Write

When any individual drive is written to, the P and Q redundancy terms must also be changed to reflect the new data (see Fig. 18). This is because the data being written over was part of a code word extending over multiple disk drives and

- 15 having P and Q terms on two disk drives. The previously stored P and Q terms will no longer be valid when part of the codeword is changed, so new P and Q terms, P'' and Q'' , must be calculated and written over the old P and Q terms on their respective disk drives. P'' and Q'' will then be proper redundancy terms for the

20 modified code word.

One possible way to calculate P'' and Q'' is to read out the whole codeword and store it in the buffers. The new portion of the codeword for drive 20C1 can then be supplied to the ACC circuit along with the rest of the codeword, and the

- 25 new P'' and Q'' can be calculated and stored on their disk drives as for a normal parallel write. However, if this method is used, it is not possible to simultaneously do another transaction mode access of a separate disk drive (i.e., drive 20A1) having part of the codeword, since that drive (20A1) and

30 its buffer are needed for the transaction mode write for the first drive (20C1).

According to a method of the present invention, two simultaneous transaction mode accesses are made possible by using only the old data to be written over and the old P and Q

- 35 to calculate the new P'' and Q'' for the new data. This is done by calculating an intermediate P' and Q' from the old data and old P and Q, and then using P' and Q' with the new data to calculate the new P'' and Q'' . This requires a read-modify-write

operation on the P and Q drives. The equations for the new P and Q redundancy are:

$$\text{New P redundancy (P'') = (old P - old data) + new data}$$

5 $\text{New Q redundancy (Q'') = (old Q - old data} \cdot a_i) + \text{new data} \cdot a_i$

$$P' = \text{old P} - \text{old data}$$

$$Q' = \text{old Q} - \text{old data} \cdot a_i$$

Where a_i is the coefficient from the syndrome equation S_1 ; and i is the index of the drive.

10

During the read portion of the read-modify-write, the data from the drive to be written to and the P and Q drives are summed by the ACC logic, as illustrated in Figure 17. This summing operation produces the P' and Q' data. The prime data is sent

15 to a data buffer. When the new data is in a data buffer, the write portion of the cycle begins as illustrated in Figure 18. During this portion of the cycle, the new data and the P' and Q' data are summed by the ACC logic to generate the new P'' and Q'' redundancy. When the summing operation is complete, the new
20 data is sent to the disk drive and the redundancy information is sent to the P and Q drives.

Parity Check of P and Q for Transaction Mode Write

During these read-modify-write operations, it is also
25 possible that the ACC unit itself may fail. In this case, if the data in a single element were to be changed by a read-modify-write operation, a hardware failure in the ACC might result in the redundancy bytes for the new data being calculated erroneously. To prevent this occurrence, the parity
30 detector and parity generator are made part of the ACC circuitry. This additional redundant circuit is shown in Figures 14a and 14b and resides within redundancy circuit 302 as shown in Fig. 11. When data is received by the ACC circuitry, parity is checked to insure that no errors have
35 occurred using the P and Q redundancy terms. In calculating Q'', new parity is generated for the product of the multiply operation and is summed with the parity of the old Q'' term. This creates the parity for the new Q term. For the P byte,

the parity bits from the data are summed with the parity bit of the old P term to create the new parity bit for the new P" term. Before writing the new data back to the disk drive, the parity of Q' (calculated as indicated previously) is checked.

- 5 Should Q' be incorrect, the second level controller engine will be informed of an ACC failure. In this manner, a failure in the ACC can be detected.

The same operations are performed for a failed disk drive write in transaction processing operations as for parallel data writes, except that data is not written to a failed drive or drives.

With respect to transaction processing functions during normal read operations, no action is required from the ACC logic. The actions taken by the ACC logic during a failed drive read in transaction processing mode are listed in Table 5 below, where i and j represent the first and second failed drives. The columns labelled Failed Drives indicate which drives have failed. The last column indicates what action the ACC may or may not take in response to the indicated failure.

20

TABLE 5

25

Failed Drives

P	-	Redundancy drives are not read; no ACC action
Q	-	Redundancy drives are not read; no ACC action
i	-	ACC logic calculates replacement data and performs a parallel read
30	i	ACC logic calculates replacement data and performs a parallel read
Q	i	ACC logic calculates replacement data and performs a parallel read
35	j	ACC logic calculates replacement data and performs a parallel read
P	Q	No ACC action as only data disk drives are read

If two data disk drives fail, the ACC logic must calculate the needed replacement data for both disk drives. If only one failed drive is to be read, both failed drives must still be noted by the ACC logic.

- 5 In the read-before-write operation (part of the read-modify-write process), the ACC logic generates P' and Q' redundancy terms. Table 6 shows the action taken by the ACC logic when a failed disk drive read precedes a write in this process. Again, i and j represent the first and second failed drives. The columns headed by Failed Drives indicate which drives have failed, and the last column denotes the response of the ACC to the indicated failures.
- 10

15

TABLE 6

Failed Drives

20	P	-	ACC calculates Q' only
	Q	-	ACC calculates P' only
	i	-	ACC logic takes no action and all good data disk drives are read into data buffers
25	i	P	All good data disk drives are read into data buffers
	Q	i	All good data disk drives are read into data buffers
	i	j	All good data disk drives are read into data buffers
30	i	failed drive	Perform a parallel read, the ACC logic calculates the replacement data for the jth failed drive. Next, the remaining good data disk drives are read into the data buffers.
	P	Q	No read before write operation is necessary

35

- When a failed data disk drive is to be written, all good data disk drives must be read so that a new P and Q redundancy can be generated. All of the data from the good data disk drive and the write data is summed to generate the new redundancy. When two data disk drives fail, the ACC logic must calculate
- 40

replacement data for both failed drives. If only one drive is to be read, both must be reported to the ACC logic.

During write operations, the ACC continues to calculate P and Q redundancy. Table 7 shows the ACC's tasks during failed drive writes. Here P and Q represent the P and Q redundancy term disk drives, and i and j represent the first and second failed data disk drives. The columns Failed Drives denote the particular failed drives, and the last column indicates the ACC response to the failed drives.

10

TABLE 7

15

Failed Drives

P	-	ACC calculates Q redundancy only
Q	-	ACC calculates P redundancy only
i	-	ACC calculates P and Q redundancy
20	i P	ACC calculates Q redundancy only
	Q i	ACC calculates P redundancy only
	i j	ACC calculates P and Q redundancy
	P Q	ACC logic takes no action

25

Summary of ECC

The interconnected arrangements herein described relative to both preferred embodiments of the present invention allow for the simultaneous transmission of data from all disks to the word assemblers or vice versa. Data from or to any given disk drive may be routed to any other word assembler through the X-bar switches under second level controller engine control. Additionally, data in any word assembler may be routed to any disk drive through the X-bar switches. The ACC units receive all data from all X-bar switches simultaneously. Any given disk drive, if it fails, can be removed from the network at any time. The X-bar switches provide alternative pathways to route data or P and Q terms around the failed component.

The parallel arrangement of disk drives and X-bar switches creates an extremely fault-tolerant system. In the prior art, a single bus feeds the data from several disk drives into a single large buffer. In the present invention, the 5 buffers are small and one buffer is assigned to each disk drive. The X-bar switches, under control of the ACC units, can route data from any given disk drive to any given buffer and vice versa. Each second level controller has several spare disks and one spare buffer coupled to it. The failure of any 10 two disks can be easily accommodated by switching the failed disk from the configuration by means of its X-bar switch and switching one of the spare disks onto the network. The present invention thus uses the error detection and correction capabilities of a Reed-Solomon error correction code in an 15 operational environment where the system's full operational capabilities can be maintained by reconfiguring the system to cope with any detected disk or buffer failure. The ACC can correct and regenerate the data for the failed disk drive and, by reconfiguring the registers of the failed and spare disk 20 drives, effectively remove the failed drive from the system and regenerate or reconstruct the data from the failed disk onto the spare disk.

Disk Drive Configuration and Format

25 The present invention allows a set of physical mass data storage devices to be dynamically configured as one or more logical mass storage devices. In accordance with the present invention, such a set of physical devices is configurable as one or more redundancy groups and each 30 redundancy group is configurable as one or more data groups.

A redundancy group, as previously used in known device sets, is a group of physical devices all of which share the same redundant device set. A redundant device is a device that stores duplicated data or check data for purposes of 35 recovering stored data if one or more of the physical devices of the group fails.

Where check data is involved, the designation of a particular physical device as a redundant device for an entire

redundancy group requires that the redundant device be accessed for all write operations involving any of the other physical devices in the group. Therefore, all write operations for the group interfere with one another, even for small data accesses 5 that involve less than all of the data storage devices.

It is known to avoid this contention problem on write operations by distributing check data throughout the redundancy group, thus forming a logical redundant device comprising portions of several or all devices of the redundancy group. 10 For example, Fig. 19 shows a group of 13 disk storage devices. The columns represent the various disks D1-D13 and the rows represent different sectors S1-S5 on the disks. Sectors containing check data are shown as hatched. Sector S1 of disk D13 contains check data for sectors of disks D1-D12. Likewise, 15 the remaining hatched sectors contain check data for their respective sector rows. Thus, if data is written to sector S4 of disk D7, then updated check data is written into sector S4 of disk D10. This is accomplished by reading the old check data, re-coding it using the new data, and writing the new 20 check data to the disk. This operation is referred to as a read-modify-write. Similarly, if data is written to sector S1 of disk D11, then check data is written into sector S1 of disk D13. Since there is no overlap in this selection of four disks 25 for writes, both read-modify-write operations can be performed in parallel.

A distribution of check data in a redundancy group in the manner shown in Fig. 19 is known as a striped check data configuration. The term "striped redundancy group" will be used herein to refer generally to a redundancy group in which 30 check data is arranged in a striped configuration as shown in Fig. 19, and the term "redundancy group stripe depth" will be used herein to refer to the depth of each check data stripe in such a striped redundancy group.

In previously known device sets, it was known to 35 provide the whole set as a single redundancy group. It has been found that a redundancy group can be divided into various "extents", each defined as a portion of the depth of the redundancy group and each capable of having a configuration of

check data different from that of other extents in the same redundancy group. Moreover, it has been found that more than one redundancy group can be provided in a single device set, under the control of a single "array controller", and connected 5 to a main processing unit via one or more device controllers.

Similarly, in previously known device sets, the single redundancy group included only one data group for application data -- i.e., the device set operated as a single logical device. It has been found, however, that a redundancy 10 group can be broken up into multiple data groups, each of which can operate as a separate logical storage device or as part of a larger logical storage device. A data group can include all available mass storage memory on a single physical device (i.e., all memory on the device available for storing 15 application data), or it can include all available mass storage memory on a plurality of physical devices in the redundancy group. Alternatively, as explained more fully below, a data group can include several physical devices, but instead of including all available mass storage memory of each device 20 might only include a portion of the available mass storage memory of each device. In addition, it has been found that it is possible to allow data groups from different redundancy groups to form a single logical device. This is accomplished, as will be more fully described, by superimposing an additional 25 logical layer on the redundancy and data groups.

Moreover, in previously known device sets in which application data is interleaved across the devices of the set, the data organization or geometry is of a very simple form. Such sets generally do not permit different logical 30 organizations of application data in the same logical unit nor do they permit dynamic mapping of the logical organization of application data in a logical unit. It has been found that the organization of data within a data group can be dynamically configured in a variety of ways. Of particular importance, it 35 has been found that the data stripe depth of a data group can be made independent of redundancy group stripe depth, and can be varied from one data group to another within a logical unit

to provide optimal performance characteristics for applications having different data storage needs.

- An embodiment of a mass storage system 500 including two second level controllers 14A and 14B is shown in the block diagram of Fig. 20. As seen in Fig. 20, each of parallel sets 501 and 502 includes thirteen physical drives 503-515 and a second level controller 14. Second level controller 14 includes a microprocessor 516a which controls how data is written and validated across the drives of the parallel set.
- Microprocessor 516a also controls the update or regeneration of data when one of the physical drives malfunctions or loses synchronization with the other physical drives of the parallel set. In accordance with the present invention, microprocessor 516a in each second level controller 14 also controls the division of parallel sets 501 and 502 into redundancy groups, data groups and application units. The redundancy groups, data groups and application units can be configured initially by the system operator when the parallel set is installed, or they can be configured at any time before use during run-time of the parallel set. Configuration can be accomplished, as described in greater detail below, by defining certain configuration parameters that are used in creating various address maps in the program memory of microprocessor 516a and, preferably, on each physical drive of the parallel set.
- Each of second level controllers 14A and 14B is connected to a pair of first level controllers 12A and 12B. Each first level controller is in turn connected by a bus or channel 522 to a CPU main memory. In general, each parallel set is attached to at least two sets of controllers so that there are at least two parallel paths from one or more CPU main memories to that parallel set. Thus, for example, each of the second level controllers 14A and 14B is connected to first level controllers 12A and 12B by buses 524 and 526. Such parallel data paths from a CPU to the parallel set are useful for routing data around a busy or failed first or second level controllers as described above.

Within each parallel set are an active set 528 comprising disk drive units 503-514, and a backup set 530

comprising disk drive unit 515. Second level controller 14 routes data between first level controllers 12 and the appropriate one or ones of disk drive units 503-515. First level controllers 12 interface parallel sets 501 and 502 to the main memories of one or more CPUs; and are responsible for processing I/O requests from applications being run by those CPUs. A further description of various components of the apparatus of parallel sets 501 and 502 and first level controllers 12 can be found in the following co-pending, commonly assigned U.S. patent applications incorporated herein in their entirety by reference: Serial No. 07/487,648 entitled "NON-VOLATILE MEMORY STORAGE OF WRITE OPERATION IDENTIFIER IN DATA STORAGE DEVICE," filed in the names of David T. Powers, Randy Katz, David H. Jaffe, Joseph S. Glider and Thomas E. Idleman; and Serial No. 07/488,750 entitled "DATA CORRECTIONS APPLICABLE TO REDUNDANT ARRAYS OF INDEPENDENT DISKS," filed in the names of David T. Powers, Joseph S. Glider and Thomas E. Idleman.

To understand how data is spread among the various physical drives of an active set 528 of a parallel set 501 or 502, it is necessary to understand the geometry of a single drive. Fig. 21 shows one side of the simplest type of disk drive -- a single platter drive. Some disk drives have a single disk-shaped "platter" on both sides of which data can be stored. In more complex drives, there may be several platters on one "spindle," which is the central post about which the platters spin.

As shown in Fig. 21, each side 600 of a disk platter is divided into geometric angles 601, of which eight are shown in Fig. 21, but of which there could be some other number. Side 600 is also divided into ring-shaped "tracks" of substantially equal width, of which seven are shown in Fig. 21. The intersection of a track and a geometric angle is known as a sector and typically is the most basic unit of storage in a disk drive system. There are fifty-six sectors 603 shown in Fig. 21.

A collection of tracks 602 of equal radius on several sides 600 of disk platters on a single spindle make up a

"cylinder." Thus, in a single-platter two-sided drive, there are cylinders of height = 2, the number of cylinders equalling the number of tracks 602 on a side 600. In a two-platter drive, then, the cylinder height would be 4. In a one-sided 5 single-platter drive, the cylinder height is 1.

A disk drive is read and written by "read/write heads" that move over the surfaces of sides 600. Fig. 22 shows the distribution of data sub-units -- sectors, tracks and cylinders -- in a group 716 of eight single-platter two-sided 10 drives 700-707 in a manner well-suited to illustrate the present invention. Drives 700-707 may, for example, correspond to drive units 503-510 of parallel set 501 or 502. Each of the small horizontal divisions represents a sector 708. For each 15 drive, four cylinders 709-712 are shown, each cylinder including two tracks 713 and 714, each track including five sectors.

In the preferred embodiment shown in Fig. 22, group 716 comprises a single redundancy group in which two types of redundancy data, referred to as "P" check data and "Q" check 20 data, are used to provide data redundancy. The P and Q check data are the results of a Reed-Solomon coding algorithm applied to the mass storage data stored within the redundancy group. The particular method of redundancy used is implementation specific. As shown, the redundancy data is distributed across 25 all spindles, or physical drives, of group 716, thus forming two logical check drives for the redundancy group comprising group 716. For example, the P and Q check data for the data in sectors 708 of cylinders 709 of drives 700-705 are contained respectively in cylinders 709 of drives 706 and 707. Each time 30 data is written to any sector 708 in any one of cylinders 709 of drives 700-705, a read-modify-write operation is performed on the P and Q check data contained in corresponding sectors of drives 706 and 707 to update the redundancy data.

Likewise, cylinders 710 of drives 700-707 share P and 35 Q check data contained in cylinders 710 of drives 704 and 705; cylinders 711 of drives 700-707 share P and Q check data contained in cylinders 711 of drives 702 and 703; and cylinders

712 of drives 700-707 share P and Q check data contained in cylinders 712 of drives 700 and 701.

Three data groups D1-D3 are shown in Fig. 22. Data group D1 includes cylinders 709 of each of spindles 700, 701. 5 Data group D2 includes cylinders 709 of each of spindles 702, 703. Data group D3 includes all remaining cylinders of spindles 700-707, with the exception of those cylinders containing P and Q check data. Data group D1 has a two-spindle bandwidth, data group D2 has a four-spindle bandwidth and data group D3 has a 10 six-spindle bandwidth. Thus it is shown in Fig. 22 that, in accordance with the principles of the present invention, a redundancy group can comprise several data groups of different bandwidths. In addition, each of data groups D1-D3 may alone, or in combination with any other data group, or groups, comprise 15 a separate logical storage device. This can be accomplished by defining each data group or combination as an individual application unit. Application units are discussed in greater detail below.

In Fig. 22, sectors 708 are numbered within each data 20 group as a sequence of logical data blocks. This sequence is defined when the data groups are configured, and can be arranged in a variety of ways. Fig. 22 presents a relatively simple arrangement in which the sectors within each of data groups D1-D3 are numbered from left to right in stripes crossing the width of the respective data group, each data 25 stripe having a depth of one sector. This arrangement permits for the given bandwidth of each data group a maximum parallel transfer rate of consecutively numbered sectors.

The term "data group stripe depth" is used herein to 30 describe, for a given data group, the number of logically contiguous data sectors stored on a drive within the boundaries of a single stripe of data in that data group. In accordance with the principles of the present invention, the depth of a data group stripe may be lesser than, greater than or equal to 35 the depth of redundancy group stripe. As one example of this, Fig. 22 shows that data groups D1-D3 each has a data group stripe depth of one sector, and are all included in a

redundancy group having a redundancy group stripe depth of one cylinder.

Redundancy group 716 can handle up to six data read requests simultaneously -- one from each of spindles 700-705 -- 5 because the read/write heads of the spindles can move independently of one another. Redundancy group 716 as configured in Fig. 22 also can handle certain combinations of write requests simultaneously. For example, in many instances any data sector of data group D1 can be written simultaneously 10 with any data sectors of data group D3 contained on spindles 702-705 that are not backed up by P or Q check data on spindles 700, 701, 706 or 707.

Redundancy group 716 as configured in Fig. 22 usually cannot handle simultaneous write operations to sectors in data 15 groups D1 and D2, however, because to perform a write operation in either of these data groups, it is necessary to write to drives 706 and 707 as well. Only one write operation can be performed on the check data of drives 706, 707 at any one time, because the read/write heads can only be in one place at one 20 time. Likewise, regardless of the distribution of data groups, write operations to any two data sectors backed up by check data on the same drive cannot be done simultaneously. The need for the read/write heads of the check drives to be in more than one place at one time can be referred to as "collision."

It is to be understood that the above-described 25 restriction concerning simultaneous writes to different data drives sharing common check drives is peculiar to check drive systems, and is not a limitation of the invention. For example, the restriction can be avoided by implementing the 30 invention using a mirrored redundancy group, which does not have the property that different data drives share redundancy data on the same drive.

Fig. 23 shows a more particularly preferred embodiment of redundancy group 716 configured according to the 35 present invention. In Fig. 23, as in Fig. 22, the logical check "drives" are spread among all of spindles 700-707 on a per-cylinder basis, although they could also be on a per-track basis or even a per-sector basis.

Data groups D1 and D2 are configured as in Fig. 22. The sectors of data group D3 of Fig. 22, however, have been divided among four data groups D4-D7. As can be seen in Fig. 23, the sequencing of sectors in data groups D4-D7 is no longer the same as the single-sector-deep striping of data groups D1 and D2. Data group D4 has a data group stripe depth of 20 sectors--equal to the depth of the data group itself. Thus, in data group D4 logically numbered sectors 0-19 can be read consecutively by accessing only a single spindle 700, thereby allowing the read/write heads of spindles 701-707 to handle other transactions. Data groups D5, D6 and D7 each show examples of different intermediate data group stripe depths of 5 sectors, 2 sectors and 4 sectors, respectively.

The distribution of the check data over the various spindles can be chosen in such a way as to minimize collisions. Further, given a particular distribution, then to the extent that second level controller 14 has a choice in the order of operations, the order can be chosen to minimize collisions.

The distribution of redundancy groups and data groups over the active set 528 of a parallel set 501 or 502 can be parameterized. For example, the redundancy group can be characterized by a redundancy group width (in spindles), representing the number of spindles spanned by a particular set of check data, a redundancy group depth (in any subunit -- sector, track or cylinder) and a redundancy group stripe depth (also in any subunit -- sector, track or cylinder). Data groups can be characterized by width (in spindles), depth (in any subunit -- sector, track or cylinder), and data group stripe depth (also in any subunit sector, track or cylinder). Because data groups do not start only at the beginning of active set 528, they are also characterized by a "base", which is a two-parameter indication of the spindle and the offset from the beginning of the spindle at which the data group starts. A redundancy group may, like a data group, include less than all of an entire spindle. In addition, as previously stated herein, a redundancy group may be divided into a plurality of extents. The extents of a redundancy group have equal widths and different bases and depths. For each extent,

the distribution of check data therein can be independently parameterized. In the preferred embodiment, each redundancy group extent has additional internal parameters, such as the depth of each redundancy group stripe within the redundancy group extent and the drive position of the P and Q check data for each such redundancy group stripe.

Redundancy group width reflects the trade-off between reliability and capacity. If the redundancy group width is high, then greater capacity is available, because only two drives out of a large number are used for check data, leaving the remaining drives for data. At the other extreme, if the redundancy group width = 4, then a situation close to mirroring or shadowing, in which 50% of the drives are used for check data, exists (although with mirroring if the correct two drives out of four fail, all data on them could be lost, while with check data any two drives could be regenerated in that situation). Thus low redundancy group widths represent greater reliability, but lower capacity per unit cost, while high redundancy group widths represent greater capacity per unit cost with lower, but still relatively high, reliability.

Data group width reflects the trade-off discussed above between bandwidth and request rate, with high data group width reflecting high bandwidth and low data group width reflecting high request rates.

Data group stripe depth also reflects a trade-off between bandwidth and request rate. This trade-off varies depending on the relationship of the average size of I/O requests to the data group and the depth of data stripes in the data group. The relationship of average I/O request size to the data group stripe depth governs how often an I/O request to the data group will span more than one read/write head within the data group; it thus also governs bandwidth and request rate. If high bandwidth is favored, the data group stripe depth is preferably chosen such that the ratio of average I/O request size to stripe depth is large. A large ratio results in I/O requests being more likely to span a plurality of data drives, such that the requested data can be accessed at a higher bandwidth than if the data were located all on one

drive. If, on the other hand, a high request rate is favored, the data group stripe depth is preferably chosen such that the ratio of I/O request size to data group stripe depth is small. A small ratio results in a lesser likelihood that an I/O 5 request will span more than one data drive, thus increasing the likelihood that multiple I/O requests to the data group can be handled simultaneously.

The variance of the average size of I/O requests might also be taken into account in choosing data group stripe 10 depth. For example, for a given average I/O request size, the data group stripe depth needed to achieve a desired request rate might increase with an increase in I/O request size variance.

In accordance with the present invention, the 15 flexibility of a mass storage apparatus comprising a plurality of physical mass storage devices can be further enhanced by grouping data groups from one or from different redundancy groups into a common logical unit, referred to herein as an application unit. Such application units can thus appear to 20 the application software of an operating system as a single logical mass storage unit combining the different operating characteristics of various data groups. Moreover, the use of such application units permits data groups and redundant groups to be configured as desired by a system operator independent of 25 any particular storage architecture expected by application software. This additional level of logical grouping, like the redundancy group and data group logical levels, is controlled by second level controller 14.

Fig. 24 illustrates an example of how application 30 units, data groups and redundancy groups might be mapped to a device set such as parallel set 501 or 502, at initialization of the parallel set.

Referring first to the linear graph 800 of logical unit address space, this graph represents the mass data storage 35 memory of the parallel set as it would appear to the application software of a CPU operating system. In the particular example of Fig. 24, the parallel set has been configured to provide a logical unit address space comprising

two application (logical) units (LUN0 and LUN1). Logical unit LUN0 is configured to include 20 addressable logical blocks having logical block numbers LBN0-LBN19. As shown by Fig. 24, logical unit LUN0 also includes an unmapped logical address 5 space 802 that is reserved for dynamic configuration. Dynamic configuration means that during run-time of the parallel set the CPU application software can request to change the configuration of the parallel set from its initial configuration. In the example of Fig. 24, unmapped spaces 802 10 and 804 are reserved respectively in each of logical units LUN0 and LUN1 to allow a data group to be added to each logical unit without requiring that either logical unit be taken off line. Such dynamic configuration capability can be implemented by providing a messaging service for a CPU application to request 15 the change in configuration. On behalf of mass storage system 500, the messaging service can be handled, for example, by the first level controllers 12. Logical unit LUN1 includes a plurality of addressable logical blocks LBN0-LBN179 and LBN200-LBN239. The logical blocks LBN180-LBN199 are reserved 20 for dynamic configuration, and in the initial configuration of the parallel set, as shown in Fig. 24, are not available to the application software.

The mass storage address space of logical unit LUN0 comprises a single data group D1, as shown by data group 25 address space chart 806. Data group D1 includes 20 logically contiguous data blocks 0-19, configured as shown in Fig. 22 and corresponding one to one with logical block numbers LBN0-LBN19. Logical unit LUN1 includes two data groups D2 and D3, comprising respectively 40 data blocks numbered 0-39 30 corresponding to logical blocks LBN200-239 of logical unit LUN1, and 180 data blocks numbered 0-179 corresponding to logical blocks LBN0-LBN179 of logical unit LUN1. As shown by the example of Fig. 24, the logical blocks of a logical unit can be mapped as desired to the data blocks of one or more data 35 groups in a variety of ways. Data group address space 806 also includes additional data groups (D4) and (D5) reserved for dynamic configuration. These data groups can be formatted on the disk drives of the parallel set at initialization or at any

time during the run-time of the parallel set, but are not available to the application software in the initial configuration of the parallel set.

The redundancy group configuration of the parallel

5 set is illustrated by a two dimensional address space 808, comprising the entire memory space of the parallel set. The horizontal axis of address space 808 represents the thirteen physical drives of the parallel set, including the twelve drives of active set 528 and the one spare drive of backup set
10 530. In Fig. 24, the drives of the active set are numbered 0-11 respectively to reflect their logical positions in the parallel set. The vertical axis of address space 808 represents the sectors of each physical drive. As shown by redundancy group address space 808, the parallel set has been
15 configured as one redundancy group RG0 having three extents A, B and C. As can be seen, the width of each extent is equal to that of the redundancy group RG0: 12 logical drive positions or, from another perspective, the entire width of active set 528.

20 Extent A of redundancy group RG0 includes sectors 1-5 of drives 0-11. Thus, extent A of redundancy group RG0 has a width of 12 spindles, and an extent depth of 5 sectors. In the example of Fig. 24, extent A is provided as memory space for diagnostic programs associated with mass storage system 500.
25 Such diagnostic programs may configure the memory space of extent A in numerous ways, depending on the particular diagnostic operation being performed. A diagnostic program may, for example, cause a portion of another extent to be reconstructed within the boundaries of extent A, including
30 application data and check data.

Extent B of redundancy group RG0 includes all application data stored on the parallel set. More particularly, in the example of Fig. 24, extent B includes data groups D1, D2 and D3 configured as shown in Fig. 22, as well as
35 additional memory space reserved for data groups (D4) and (D5), and a region 809 of memory space not mapped to either logical unit LUN0 or LUN1. This region 809 may, for example, be mapped

to another logical unit (e.g., LUN2) being used by another application.

Address space 808 also includes a third extent C in which a second diagnostic field may be located. Although the 5 parallel set is shown as including only a single redundancy group RG0, the parallel set may alternatively be divided into more than one redundancy group. For example, redundancy group RG0 might be limited to a width of 8 spindles including logical drive positions 0-7, such as is shown in Figs. 22 and 23, and a 10 second redundancy group might be provided for logical drive positions 8-11.

It is also not necessary that the entire depth of the parallel set be included in redundancy group RG0. As an example, Fig. 24 shows that above and below redundancy group 15 RG0 are portions 810 and 811 of memory space 808 that are not included in the redundancy group. In the example of Fig. 24, portions 810 and 811 contain data structures reflecting the configuration of the parallel set. These data structures are described in greater detail below in connection with Fig. 25. 20 In addition, any portion of memory space between set extents A, B and C, such as the portions indicated by regions D and E in Fig. 24, may be excluded from redundancy group RG0.

Fig. 24 further provides a graph 812 showing a linear representation of the physical address space of the drive in 25 logical position 0. Graph 812 represents a sectional view of address space chart 810 along line 0'-0", and further illustrates the relationship of the various logical levels of the present invention as embodied in the exemplary parallel set configuration of Fig. 24.

30 As stated previously, the parallel set can be configured by the operator initially at installation time and/or during run-time of the parallel set. The operator formats and configures the application units he desires to use by first determining the capacity, performance and redundancy 35 requirements for each unit. These considerations have been previously discussed herein. Once the capacity, performance and redundancy requirements have been defined, the logical structure of the units can be specified by defining parameters

for each of the logical layers (redundancy group layer, data group layer and application unit layer). These parameters are provided to a configuration utility program executed by processor 516a of second level controller 14. The 5 configuration utility manages a memory resident database of configuration information for the parallel set. Preferably, a copy of this database information is kept in non-volatile memory to prevent the information from being lost in the event of a power failure affecting the parallel set. A format 10 utility program executed by processor 516a utilizes the information in this database as input parameters when formatting the physical drives of the parallel set as directed by the operator.

The basic parameters defined by the configuration 15 database preferably include the following:

1) For each redundancy group:

- 20 Type: Mirrored;
Two check drives;
One check drive;
No check drive.
- 25 Width: The number of logical drive positions as spindles in the redundancy group.
- 30 Extent Size: For each extent of the redundancy group, the size (depth) of the extent in sectors.
- Extent Base: For each extent of the redundancy group the physical layer address of the first sector in the extent.
- 35 Stripe in Depth: For interleaved check drive groups, the depth, sectors of a stripe of check data.
- Drives: An identification of the physical drives included in the redundancy group.
- 40 Name: Each redundancy group has a name that is unique across the mass storage system 500.
- 45 Base: The index (logical drive number) of the drive position within the redundancy group that is the first drive position in the data group within the redundancy group.

	Width:	The number of drive positions (logical drives) in the data group. This is the number of sectors across in the data group address space.
5	Start:	The offset, in sectors, within the redundancy group extent where the data group rectangle begins on the logical drive position identified by the base parameter.
10	Depth:	The number of sectors in a vertical column of the data group, within the redundancy group extent. Depth and width together are the dimensions respectively of the side and the top of the rectangle formed by each data group as shown in Figs. 22-24.
15	Redundancy Group:	The name of the redundancy group to which the data group belongs.
20	Extent Number:	A name or number identifying the extent in which the data group is located.
25	Index:	The configuration utility will assign a number to each data group, unique within its redundancy group. This number will be used to identify the data group later, for the format utility and at run-time.
30	Data Group Stripe Depth:	The depth, in sectors, of logically contiguous blocks of data within each stripe of data in the data group.
35	3) For each application unit:	
	Size:	Size in sectors
40	Data Group List:	A list of the data groups, and their size and order, within the unit address space, and the base unit logical address of each data group. Each group is identified by the name of the redundancy group it is in and its index.
45	Fig. 25 illustrates exemplary data structures containing the above-described parameters that can be used in implementing the configuration database of a device set such as parallel set 501 or 502. These data structures may be varied as desired to suit the particular device set embodiment to which they are applied. For example, the data structures described hereafter allow for many options that may be unused	
50		

in a particular device set, in which case the data structures may be simplified.

The configuration database includes an individual unit control block (UCB) for each application unit that

- 5 references the parallel set (a unit may map into more than one parallel set). These UCB's are joined together in a linked list 900. Each UCB includes a field labeled APPLICATION UNIT # identifying the number of the application unit described by that UCB. Alternatively, the UCB's within link list 900 might
- 10 be identified by a table of address pointers contained in link list 900 or in some other data structure in the program memory of microprocessor 516a. Each UCB further includes a map 901 of the data groups that are included in that particular application unit. Data group map 901 includes a count field
- 15 902 defining the number of data groups within the application unit, a size field 904 defining the size of the application unit in sectors, and a type field 906 that defines whether the linear address space of the application unit is continuous (relative addressing) or non-continuous (absolute addressing).
- 20 A non-continuous address space is used to allow portions of the application unit to be reserved for dynamic configuration as previously described in connection with data groups (D4) and (D5) of Fig. 22.

Data group map 901 further includes a data group mapping element 908 for each data group within the application unit. Each data group mapping element 908 includes a size field 910 defining the size in sectors of the corresponding data group, a pointer 912 to a descriptor block 914 within a data group list 916, a pointer 718 to an array control block 720, and an index field 721. The data group mapping elements 908 are listed in the order in which the data blocks of each data group map to the LBN's of the application unit. For example, referring to LUN1 of Fig. 24, the mapping element for data group D3 would be listed before the data group mapping element for data group D2. Where the address space of the application unit is non-continuous, as in the case of LUN1 of Fig. 24, data group map 901 may include mapping elements

corresponding to, and identifying the size of, the gaps between available ranges of LBN's.

Data group list 916 includes a descriptor block 914 for each data group within the parallel set, and provides 5 parameters for mapping each data group to the redundancy group and redundancy group extent in which it is located. Data group list 916 includes a count field 717 identifying the number of descriptor blocks in the list. In the case of a redundancy group having a striped check data configuration, each data 10 group descriptor block 914 may include a "pqdel" field 722 that defines the offset of the first data block of the data group from the beginning of the check data for the redundancy group stripe that includes that first data block. The value of pqdel field 722 may be positive or negative, depending on the 15 relative positions of the drive on which the first data block of the data group is configured and the corresponding check data drives for the redundancy group stripe including that first data block. This value can be useful for assisting the second level controller in determining the position of the 20 check data during I/O operations.

Each data group descriptor block 914 also includes an index field 723 (same value as index field 721), a width field 724, a base field 726, an extent number field 727, a start field 728, a depth field 730, a data group stripe depth field 25 731 and a redundancy group name field 732 that respectively define values for the corresponding parameters previously discussed herein.

Array control block 720 provides a map of redundancy groups of the parallel set to the physical address space of the 30 drives comprising the parallel set. Array control block 720 includes an array name field 734 and one or more fields 735 that uniquely identify the present configuration of the parallel set. Array control block 720 also includes a list of redundancy group descriptor blocks 736. Each redundancy group 35 descriptor block 736 includes a redundancy group name field 738 identifying the redundancy group corresponding to the descriptor block, a redundancy group width field 740 and a redundancy group extent map 742. Array control block 720

further includes a list 744 of physical drive identifier blocks 745.

- For each extent within the redundancy group, redundancy group extent map 742 includes an extent descriptor block 746 containing parameters that map the extent to corresponding physical address in the memory space of the parallel set, and define the configuration of redundant information in the extent. As an example, extent descriptor blocks are shown for the three extents of redundancy group RG0 of Fig. 24, each extent descriptor block including an extent number field 747 and base and size fields defining the physical addresses of the corresponding extent. Application data base and size fields 748 and 750 correspond respectively to the base and size of extent B of redundancy group RG0; diagnostic (low) base and size fields 752 and 754 correspond respectively to the base and size of extent A of redundancy group RG0; and diagnostic (high) base and size fields 756 and 758 correspond respectively to the base and size of extent C of redundancy group RG0.
- Each extent descriptor block 746 also includes a type field 760 that defines the type of redundancy implemented in the extent. For example, a redundancy group extent may be implemented by mirroring or shadowing the mass storage data stored in the data group(s) within the extent (in which case, the extent will have an equal number of data drives and redundant drives). Alternatively, a Reed-Solomon coding algorithm may be used to generate check data on one drive for each redundancy group stripe within the extent, or a more sophisticated Reed-Solomon coding algorithm may be used to generate two drives of check data for each redundancy group stripe. Type field 760 may specify also whether the check data is to be striped throughout the extent, and how it is to be staggered (e.g., the type field might index a series of standardized check data patterns, such as a pattern in which check data for the first redundancy group stripe in the extent is located on the two numerically highest logical drive positions of the redundancy group, check data for the second redundancy group stripe in the extent is located on the next

two numerically highest logical drive positions, and so on). Yet another alternative is that type field 760 indicates that no check drives are included in the initial configuration of the redundancy group extent. This may be desired, for example,

5 if the redundancy group extent is created for use by diagnostic programs. A redundancy group extent of this type was previously discussed in connection with extent A of redundancy group RG0 shown in Fig. 24.

Each extent descriptor block 746 may further include

10 a redundancy group stripe depth field 762 to specify, if appropriate, the depth of redundancy group stripes within the extent.

List 744 of physical drive identifier blocks 745 includes an identifier block 745 for each physical drive in the parallel set. Each identifier block 745 provides information concerning the physical drive and its present operating state, and includes in particular one or more fields 764 for defining the logical position in the parallel set of the corresponding physical drive.

20 To summarize briefly the intended functions of the various data structures of Fig. 25, the unit control blocks of link list 900 define the mapping of application units to data groups within the parallel set. Mapping of data groups to redundancy groups is defined by data group list 916, and

25 mapping of redundancy groups to the physical address space of the memory of the parallel set is defined by array control block 720.

When each physical disk of the parallel set is formatted by the formatting utility, a copy of the array

30 control block 720, link list 900 and data group list 916 are stored on the drive. This information may be useful for various operations such as reconstruction of a failed drive. A copy of the configuration database also may be written to the controller of another parallel set, such that if one parallel

35 set should fail, another would be prepared to take its place.

During each I/O request to a parallel set, the mapping from unit address to physical address spaces must be made. Mapping is a matter of examining the configuration

- database to translate: (1) from a unit logical address span specified in the I/O request to a sequence of data group address spans; (2) from the sequence of data group address spans to a set of address spans on logical drive positions
- 5 within a redundancy group; and then (3) from the set of address spans on logical drive positions to actual physical drive address spans. This mapping process can be done by having an I/O request server step through the data structures of the configuration database in response to each I/O request.
- 10 Alternatively, during initialization of the parallel set the configuration utility may, in addition to generating the configuration database as previously described, generate subroutines for the I/O request server for performing a fast mapping function unique to each data group.' The particular
- 15 manner in which the I/O request server carries out the mapping operations is implementation specific, and it is believed to be within the skill of one in the art to implement an I/O request server in accordance with the present invention as the invention is described herein.
- 20 The following is an example of how the I/O request server might use the data structures of Fig. 25 to map from a logical unit address span of an application I/O request to a span or spans within the physical address space of a parallel set. The logical unit address span is assumed to be defined in
- 25 the I/O request by a logical application unit number and one or more LBN's within that application unit.
- The I/O request server determines from the I/O request the application unit being addressed and whether that application unit references the parallel set. This latter
- 30 determination can be made by examining link list 900 for a UCB having an APPLICATION UNIT # corresponding to that of the I/O request. If an appropriate UCB is located, the I/O request server next determines from the LBN(s) specified in the I/O request the data group or data groups in which data block(s)
- 35 corresponding to those LBN(s) are located. This can be accomplished by comparing the LBN(s) to the size fields 910 of the mapping elements in data group map 901, taking into account the offset of that size field from the beginning of the

application unit address space (including any gaps in the application unit address space). For example, if the size value of the first data group mapping element in map 901 is greater than the LBN(s) of the I/O request, then it is known 5 that the LBN(s) correspond to data blocks in that data group. If not, then the size value of that first mapping element is added to the size value of the next mapping element in map 901 and the LBN(s) are checked against the resulting sum. This process is repeated until a data group is identified for each 10 LBN in the I/O request.

Having identified the appropriate data group(s), the I/O request server translates the span of LBN's in the I/O request into one or more spans of corresponding data block numbers within the identified data group(s). The configuration 15 utility can then use the value of index field 921 and pointer 912 within each mapping element 908 corresponding to an identified data group to locate the data group descriptor block 914 in data group list 916 for that data group. The I/O request server uses the parameters of the data group descriptor block 20 to translate each span of data block numbers into a span of logical drive addresses.

First, the I/O request server determines the logical drive position of the beginning of the data group from the base field 726 of the data group descriptor block 914. The I/O 25 request server also determines from fields 732 and 727 the redundancy group name and extent number in which the data group is located, and further determines from start field 728 the number of sectors on the drive identified in base field 726 between the beginning of that redundancy group extent and the 30 beginning of the data group. Thus, for example, if the I/O request server is reading the descriptor block for data group D3 configured as shown in Fig. 24, base field 726 will indicate that the data group begins on logical drive position 0, redundancy name field 732 will indicate that the data group is 35 in redundancy group RG0, extent field 727 will indicate that the data group is in extent B, and start field 728 will indicate that there is an offset of 10 sectors on logical drive

0 between the beginning of extent B and the first data block of data group D3.

- Knowing the logical drive position and extent offset of the first data block of the data group, the I/O request server then determines the logical drive position and extent offset for each sequence of data blocks in the data group corresponding to the LBN's of the I/O request. To do this, the I/O request server may use the values of width field 724, depth field 730 and data group stripe depth field 731. If any check data is included within the rectangular boundaries of the data group, the position of the check data is taken into account if necessary in determining the logical drive position and extent offset address spans of the data blocks. This can be accomplished using information from array control block 720.
- More particularly, the I/O request server can determine the logical drive position and extent offset of any check data within the boundaries of the data group by examining the type field 760 and the redundancy group stripe depth field 762 of the appropriate redundancy group extent descriptor block 746 (the I/O request server can determine which extent descriptor block 746 is appropriate by finding the extent descriptor block 746 having an extent number field 747 that matches the corresponding extent number field 727 in the data group's descriptor block 914). The I/O request server is directed to array control block 720 by the pointer 718 in the data group mapping element 908.

- To translate each logical drive position and extent offset address span to a physical address span on a particular physical drive of the parallel set, the I/O request server reads the physical drive identifier blocks 745 to determine the physical drive corresponding to the identified logical drive position. The I/O request server also reads the base field of the appropriate extent descriptor block 746 of array control block 720 (e.g., application base field 752), which provides the physical address on the drive of the beginning of the extent. Using the extent offset address span previously determined, the I/O request server can then determine for each

physical drive the span of physical addresses that corresponds to the identified extent offset address span.

It may occur that during operation of a parallel set one or more of the physical drives is removed or fails, such

- 5 that the data on the missing or failed drive must be reconstructed on a spare drive. In this circumstance, the configuration of the set must be changed to account for the new drive, as well as to account for temporary set changes that must be implemented for the reconstruction period during which
- 10 data is regenerated from the missing or failed drive and reconstructed on the spare. It is noted that the configuration utility can be used to remap the set configuration by redefining the parameters of the configuration database.

In general, to those skilled in the art to which this invention relates, many changes in construction and widely differing embodiments and applications of the present invention will suggest themselves without departing from its spirit and scope. For instance, a greater number of second level controllers and first level controllers may be implemented in the system. Further, the structure of the switching circuitry connecting the second level controllers to the disk drives may be altered so that different drives are the primary responsibility of different second level controllers. Thus, the disclosures and descriptions herein are purely illustrative and not intended to be in any sense limiting. The scope of the invention is set forth in the appended claims.

WHAT IS CLAIMED IS:

1. A system for storing data received from an external source, comprising:

- 5 at least two control means for providing control of data flow to and from the external source;
- 10 a plurality of storage means coupled to at least two control means wherein said storage means are divided into groups and each group is controlled by at least two of said control means such that in the case that a first control means coupled to a particular group of storage means fails, control of said particular group is assumed by a second control means;
- 15 a plurality of data handling means coupled to at least two control means for disassembling data into data blocks to be written across a group of said storage means; and
- 20 error detection means coupled to said control means and said storage means for calculating at least one error detection term for each group of storage means based on the data received from the external source using a selected error code and providing said error detection term to be compared with data to detect errors, said error detection means being coupled to each of said control means to receive the data from said control means and transmit said error detection term to an error code storage means in said group of storage means.

25

2. The system of claim 1 wherein said data handling means further includes assembly means for assembling said data blocks received from said control means.

30

3. The system of claim 1 further comprising a plurality of buffer means coupled to a first bus and to said control means for buffering data received by and transmitted from the system.

35

4. The system of claim 3 wherein said control means further comprises:

 a plurality of switching means, each switching means being coupled both to each storage means and each buffer means

for providing switchable control between any of said storage means and any of said buffer means; and

5 switch control means coupled to each switching means for controlling said switching means to allow data to flow from a selected buffer means to a selected storage means and to flow from said selected storage means to said selected buffer means.

10 5. The system of claim 3 further comprising error correction means coupled to said error detection means for correcting error in data as said data is transmitted from either said buffer means to said storage means through said switching means or from said storage means to said buffer means through said switching means.

15 6. The system of claim 3 wherein the error detection means uses a Reed-Solomon error code to detect errors in said data received from said buffer means and said storage means.

20 7. The system of claim 3 further comprising detachment means coupled to said error detection means for detaching from the system storage means and buffer means which transmit erroneous data responsive to receiving said error detection term from said error detection means.

25 8. The system of claim 3 wherein said error detection means routes data from any selected buffer means to any selected storage means and from any selected storage means to any selected buffer means.

30 9. The system of claim 1 wherein said plurality of storage means comprises:

a first group of data storage means for storing data from the external source; and

35 a second group of error check and correction (ECC) storage means for storing ECC data generated by said error detection means.

10. The system of claim 1 wherein each of said plurality of storage means stores data and error check and correction (ECC) data in a predefined pattern.

5 11. The system of claim 1 wherein each of said plurality of control means comprises:

a plurality of first-level means for handling input/output storage communication with the external source; and

10 a plurality of second-level means, each connected to each of said plurality of first-level means and at least one partner second-level means, for providing a data path from said first-level means to each of a primary group of said storage means, and for providing a data path to a secondary group of
15 said storage means.

12. The system of claim 11 wherein each of said plurality of second-level means is configured to maintain a particular logical configuration of said plurality of storage means transparent to the external source in case a failure occurs in said partner second-level means by providing a data path to a secondary group of said storage means formerly primarily controlled by said at least one partner second-level means.

25

13. The system of claim 11 further comprising first communication lines connected between each of said plurality of second-level means such that when one of said second-level means fails, said one second-level means informs an other 30 second-level means that a failure has occurred so that said other second-level means can assume primary control of a subset of storage means formerly primarily controlled by said one second-level means.

35

14. The system of claim 11 further comprising a switching function connected between said plurality of second-level means and said plurality of storage means for

transferring control from a first second-level means upon failure to a second second-level means.

5 15. The system of claim 11 wherein each second-level means further comprises a second-level means recovery system for issuing signals to and receiving signals from said first-level means and said at least one partner second-level means.

10 16. The system of claim 11 wherein each second-level means further comprises a state machine capable of maintaining said second-level controller in a number of states, each state representing a system configuration wherein said second-level means controls particular storage means.

15 17. The system of claim 1 wherein said plurality of storage means are operatively interconnected to function at a first logical level as a plurality of redundancy groups, each of said redundancy groups including at a second logical level at least one data group, each data group capable of operating 20 as a separate logical storage device.

18. The system of claim 17 wherein each redundancy group comprises a plurality of data groups.

25 19. The system of claim 17 further comprising a third logical level, wherein at least one data group from each of at least two of the plurality of redundancy groups are combined to form a single logical mass data storage device.

30 20. The system of claim 17 wherein for at least one of the plurality of redundancy groups redundancy is provided by an error detecting and correcting code, the code words of which are stored in at least one check drive included in at least one redundancy group.

35

21. The system of claim 20 wherein each of the at least one check drives is a particular storage means.

22. The system of claim 20 wherein each of the at least one check drive is a logical mass storage device comprising portions of a plurality of said storage means.

5 23. The system of claim 17 wherein for at least one of the plurality of redundancy groups redundancy is provided by mirroring.

10 24. The system of claim 17 wherein redundancy data is stored in at least one redundant mass storage means included in the redundancy group.

15 25. In a system including at least two control means for communicating with an external source and a plurality of storage means wherein at least two of the control means are connected to each of the storage means, a method for storing data received from the external source comprising the steps of:
receiving data from the external source;
configuring the plurality of storage means into
20 groups wherein each group is initially controlled by at least two of the control means such that in the case that one of the control means fails, the storage means of each group is accessible through another one of the control means;
disassembling data into groups of data blocks to be
25 written to said plurality of storage means;
calculating at least one error detection term from said data using a selected error code;
storing said data blocks in a first of said groups of storage means; and
30 storing said at least one error detection term in said first of said groups of storage means.

26. The method of claim 25 further comprising the steps of:

35 retrieving said data blocks from said first of said groups of storage means;
calculating a check error detection term from said data blocks using a selected error code;

retrieving said at least one error detection term from said first of said groups of storage means; and

comparing said check error detection term to said at least one error detection term to determine that said data has not been corrupted.

5

27. The method of claim 26 further comprising the step of correcting said data if it is determined that said data has been corrupted.

10

28. The method of claim 26 further comprising the step of assembling said data blocks into a form in which it was received from the external source.

15

29. The method of claim 25 wherein the step of configuring further comprises the step of setting a plurality of switching means to allow said data blocks to be passed between the control means and the storage means in a predefined pattern.

20

30. The method of claim 25 further comprising the step of detaching a particular storage means upon which said data was stored if it is determined that said data has been corrupted.

25

31. The method of claim 25 further comprising the steps of:

detecting a failure in a first second-level controller in said control means; and

30

switching a data path from said first second-level controller to a second second-level controller transparently to the external source such that the storage means formerly controlled by said first second-level controller are maintained in communication with the external source.

35

32. The method of claim 25 further comprising the step of configuring said storage means to function at a first logical level as a plurality of redundancy groups, each of said

redundancy groups including at a second logical level at least one data group wherein each data group is capable of operating as a separate logical storage device.

5 33. The method of claim 32 wherein each redundancy group comprises a plurality of data groups.

10 34. The method of claim 32 further comprising the step of configuring the storage means with a third logical level, wherein at least one data group from each of at least two of the plurality of redundancy groups are combined to form a single logical mass storage device.

15 35. The method of claim 32 wherein for at least one of the plurality of redundancy groups redundancy is provided by an error detecting and correcting code, the code words of which are stored in at least one check storage means included in at least one redundancy group.

20 36. The method of claim 35 wherein each of the at least one check storage means is a particular storage means.

25 37. The method of claim 35 wherein each of the at least one check storage means is a logical mass storage device comprising portions of a plurality of the storage means.

30 38. The method of claim 35 wherein for at least one of the plurality of redundancy groups redundancy is provided by mirroring.

35 39. The method of claim 35 wherein redundancy data is stored in at least one redundant storage means included in the redundancy group.

40. A system for storing data received from an external source, comprising:

control means for providing control of data flow to and from the external source;

a plurality of storage means coupled to said control means wherein said storage means are divided into groups;

a plurality of data handling means coupled to said control means for disassembling data with data blocks to be written to said storage means; and

error detection means coupled to said control means for receiving said data blocks in parallel form and detecting errors in each data block substantially simultaneously as said data blocks are written to said storage means.

10

41. The system of claim 40 further comprising data correction means coupled to said error detection means for correcting corrupted data in response to an error detection signal provided by said error detection means.

15

42. The system of claim 41 wherein the error detection means uses a Reed-Solomon error code to detect errors in said data received from said buffer means and said storage means.

20

43. The system of claim 41 wherein the error correction means uses a Reed-Solomon error code to correct errors in said data received from said buffer means and said storage means.

25

44. The system of claim 40 further comprising detachment means coupled to said error detection means for detaching a particular storage means from the system which has provided corrupted data as determined by said error detection means.

30

45. The system of claim 40 wherein said data handling means further includes assembly means for assembling said data blocks received from said control means.

35

46. The system of claim 40 further comprising a plurality of buffer means coupled to a first bus and to said control means for buffering data received by and transmitted from the system.

47. The system of claim 40 wherein said control means further comprises:

5 a plurality of switching means, each switching means being coupled both to each storage means and each buffer means for providing switchable control of between any of said storage means and any of said buffer means; and

10 switch control means coupled to each switching means for controlling said switching means to allow data to flow from a selected buffer means to a selected storage means and to flow from said selected storage means to said selected buffer means.

15 48. The system of claim 47 wherein said error detection means routes data from any selected buffer means to any selected storage means and from any selected storage means to any selected buffer means.

20 49. The system of claim 41 wherein said plurality of storage means comprises:

a first group of data storage means for storing data from the external source; and

25 a second group of error check and correction (ECC) storage means for storing ECC data generated by said error correction means.

30 50. The system of claim 40 wherein each of said plurality of storage means stores data and error check and correction (ECC) data in a predefined pattern.

35 51. The system of claim 40 wherein said control means comprises:

a plurality of first-level means for handling input/output storage communication with the external source; and

35 a plurality of second-level means, each connected to each of said plurality of first-level means and at least one partner second-level means, for providing a data path from said first-level means to each of a primary group of said storage

means, and for providing a data path to a secondary group of said storage means.

52. The system of claim 51 wherein each of said plurality of second-level means is configured to maintain a particular logical configuration of said plurality of storage means transparent to the external source in case a failure occurs in said partner second-level means by providing a data path to a secondary group of said storage means formerly primarily controlled by said at least one partner second-level means.

53. The system of claim 51 further comprising first communication lines connected between each of said plurality of second-level means such that when one of said second-level means fails, said one second-level means informs an other second-level means that a failure has occurred so that said other second-level means can assume primary control of a subset of storage means formerly primarily controlled by said one second-level means.

54. The system of claim 51 further comprising a switching function connected between said plurality of second-level means and said plurality of storage means for transferring control from a first second-level means upon failure to a second second-level means.

55. The system of claim 51 wherein each second-level means further comprises a second-level means recovery system for issuing signals to and receiving signals from said first-level means and said at least one partner second-level means.

56. The system of claim 51 wherein each second-level means further comprises a state machine capable of maintaining said second-level controller in a number of states, each state representing a system configuration wherein said second-level means controls particular storage means.

57. The system of claim 40 wherein said plurality of storage means are operatively interconnected to function at a first logical level as a plurality of redundancy groups, each of said redundancy groups including at a second logical level 5 at least one data group, each data group capable of operating as a separate logical storage device.

58. The system of claim 57 wherein each redundancy group comprises a plurality of data groups.

10

59. The system of claim 57 further comprising a third logical level, wherein at least one data group from each of at least two of the plurality of redundancy groups are combined to form a single logical mass data storage device.

15

60. The system of claim 57 wherein for at least one of the plurality of redundancy groups redundancy is provided by an error detecting and correcting code, the code words of which are stored in at least one check drive included in at least one 20 redundancy group.

61. The system of claim 60 wherein each of the at least one check drives is a particular storage means.

25

62. The system of claim 60 wherein each of the at least one check drive is a logical mass storage device comprising portions of a plurality of said storage means.

30

63. The system of claim 57 wherein for at least one of the plurality of redundancy groups redundancy is provided by mirroring.

35

64. The system of claim 57 wherein redundancy data is stored in at least one redundant mass storage means included in the redundancy group.

65. In a system including control means for communicating with an external source and a plurality of

storage means, a method for storing data received from the external source comprising the steps of:

receiving data from the external source;

disassembling the data into groups of data blocks to

5 be written to said plurality of storage means;

calculating at least one error detection for each data block substantially simultaneously; and

storing said data blocks and at least one error detection term in a first of said groups of storage means

10 substantially simultaneously.

66. The method of claim 65 further comprising the steps of:

15 retrieving said data blocks from said first of said groups of storage means;

calculating a check error detection term from said data blocks using a selected error code;

retrieving said at least one error detection term from said first of said groups of storage means; and

20 comparing said check error detection term to said at least one error detection term to determine that said data has not been corrupted.

67. The method of claim 66 further comprising the step of correcting said data if it is determined that said data has been corrupted.

68. The method of claim 66 further comprising the step of assembling said data blocks into a form in which it was received from the external source.

69. The method of claim 65 wherein the step of configuring further comprises the step of setting a plurality of switching means to allow said data blocks to be passed between the control means and the storage means in a predefined pattern.

70. The method of claim 65 further comprising the step of detaching a particular storage means upon which said data was stored if it is determined that said data has been corrupted.

5

71. The method of claim 65 further comprising the steps of:

detecting a failure in a first second-level controller in said control means; and

10 switching a data path from said first second-level controller to a second second-level controller transparently to the external source such that the storage means formerly controlled by said first second-level controller are maintained in communication with the external source.

15

72. The method of claim 65 further comprising the step of configuring said storage means to function at a first logical level as a plurality of redundancy groups, each of said redundancy groups including at a second logical level at least 20 one data group wherein each data group is capable of operating as a separate logical storage device.

73. The method of claim 72 wherein each redundancy group comprises a plurality of data groups.

25

74. The method of claim 72 further comprising the step of configuring the storage means with a third logical level, wherein at least one data group from each of at least two of the plurality of redundancy groups are combined to form 30 a single logical mass storage device.

75. The method of claim 72 wherein for at least one of the plurality of redundancy groups redundancy is provided by an error detecting and correcting code, the code words of which 35 are stored in at least one check storage means included in at least one redundancy group.



76. The method of claim 75 wherein each of the at least one check storage means is a particular storage means.

77. The method of claim 75 wherein each of the at least one check storage means is a logical mass storage device comprising portions of a plurality of the storage means.

78. The method of claim 75 wherein for at least one of the plurality of redundancy groups redundancy is provided by mirroring.

79. The method of claim 75 wherein redundancy data is stored in at least one redundant storage means included in the redundancy group.

1/31

**FIG. 1.** (PRIOR ART)**SUBSTITUTE SHEET****FIG. 2.** (PRIOR ART)

2/31




FIG. 3.

3/31




FIG. 3.5

SUBSTITUTE SHEET

4/31






FIG.-4.

SUBSTITUTE SHEET

5/31

**SUBSTITUTE SHEET****FIG. - 5.**



SUBSTITUTE SHEET

7/31





FIG.—7.

8/31



FIG_8A.

SUBSTITUTE SHEET




FIG.-8B.




FIG.-8E.

SUBSTITUTE SHEET

10/31




FIG.—8C.

11/31




FIG.-8D.

SUBSTITUTE SHEET

12/31




FIG.—8F.

SUBSTITUTE SHEET

13/31







FIG.—8G.



SUBSTITUTE SHEET

15/31

**FIG. 9.****SUBSTITUTE SHEET**

16/31




FIG. 10.

17/31




FIG. II.

SUBSTITUTE SHEET

18/31




FIG. 12A.




FIG. 12B SUBSTITUTE SHEET

19/31

CODE WORD A		400			
d_{N-1}	d_{N-2}	\dots	d_1	d_0	P Q




FIG._/3A.

CODE WORD		450			
d_{N-1}	d_{N-1}	\dots	d_1	d_0	P Q

CODE WORD		SECTOR A P DRIVE			
d_{N-1}	d_{N-1}	\dots	d_1	d_0	P Q

CODE WORD		SECTOR A Q DRIVE			
d_{N-1}	d_{N-1}	\dots	d_1	d_0	P Q

FIG._/3B.

SUBSTITUTE SHEET

20/31




FIG. 14A.




FIG. 14B.

SUBSTITUTE SHEET

22/31




FIG. 16.

21/31

TRANSACTION MODE:
NORMAL READ




FIG. 15.

SUBSTITUTE SHEET

24/31

TRANSACTION MODE:
READ-MODIFY-WRITE:
WRITE




FIG.-18.

SUBSTITUTE SHEET

23/31

TRANSACTION MODE:
READ-MODIFY-WRITE:
READ




FIG.—17.

SUBSTITUTE SHEET

26/31




FIG. 20.

SUBSTITUTE SHEET

25/31




FIG.—19.

SUBSTITUTE SHEET

28/31

700	701	D ₁	702	703	704	705	706	707
708	0	1	0	5	2	3		
716	2	3	4	9	6	7		
	4	5	11	13	10	11		
	6	7	12	14	13	15		
	8	9	16	17	18	19		
D ₁	10	11	20	21	22	23	P	
	12	13	24	25	26	27		Q
	14	15	28	29	30	31		
	16	17	32	33	34	35		
	18	19	36	37	38	39		
	0	1	2	3			4	5
	6	7	8	9			10	11
D ₃	12	13	14	15			16	17
	18	19	20	21			22	23
	24	25	26	27			28	29
	30	31	32	33			34	35
	36	37	38	39			40	41
	42	43	44	45			46	47
	48	49	50	51			52	53
	54	55	56	57			58	59
	60	61			62	63	64	65
	66	67			68	69	70	71
	72	73			74	75	76	77
	78	79			80	81	82	83
	84	85			86	87	88	89
	90	91		Q	92	93	94	95
	96	97			98	99	100	101
	102	103			104	105	106	107
	108	109			110	111	112	113
	114	115			116	117	118	119
			120	121	122	123	124	125
			126	127	128	129	130	131
			132	133	134	135	136	137
			138	139	140	141	142	143
			144	145	146	147	148	149
			150	151	152	153	154	155
			156	157	158	159	160	161
			162	163	164	165	166	167
			168	169	170	171	172	173
			174	175	176	177	178	179
							713	714

FIG. 22.

SUBSTITUTE SHEET

27/31




FIG. 21.

SUBSTITUTE SHEET

29/31

716

D₂

0	1
2	3
4	5
6	7
8	9
10	11
12	13
14	15
16	17
18	19
0	20
1	21
2	22
3	23
4	24
5	25
6	26
7	27
8	28
9	29
10	30
11	31
12	32
13	33
14	34
15	35
16	36
17	37
18	38
19	39

D₁D₄


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39																
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39																
P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q	P	Q																				
40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79																
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39																
24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79

RED. GROUP

PHYSICAL GROUP

FIG. 23.

SUBSTITUTE SHEET

**SUBSTITUTE SHEET**




FIG. 25.

INTERNATIONAL SEARCH REPORT

International Application No. PCT/US91/01276

I. CLASSIFICATION OF SUBJECT MATTER (If several classification symbols apply, indicate all) *

According to International Patent Classification (IPC) or to both National Classification and IPC

IPC(5): G06F 7/22, 11/10
U.S. CL.: 371/8.2,11.2; 364/245.3

II. FIELDS SEARCHED

Minimum Documentation Searched ?

Classification System	Classification Symbols
U.S.	371/8.2,10.1,11.1,11.2 364/245.3,245.4

Documentation Searched other than Minimum Documentation
to the Extent that such Documents are Included in the Fields Searched *

III. DOCUMENTS CONSIDERED TO BE RELEVANT *

Category *	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³
P, X Y	US, A, 4,914,656 (DUNPHY, JR. ET AL.) 03 April 1990. See column 2, line 36 through column 4, line 24; column 5, line 62 through column 6, line 14.	1,2,9, 10,25- 30,40, 41,44, 45,49, 50,65- 70 3,5,6, 7,42, 43,46
A	US, A, 4,849,929 (TIMSIT) 18 July 1989. See the entire document.	1-79
A	US, A, 4,825,403 (GERSHENSON ET AL.) 25 April 1989. See the entire document.	1-79
A	US, A, 4,722,085 (FLORA ET AL.) 26 January 1988. See the entire document.	1-79

* Special categories of cited documents: ¹⁰

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

IV. CERTIFICATION

Date of the Actual Completion of the International Search

15 MAY 1991

Date of Mailing of this International Search Report

13 JUN 1991

International Searching Authority

ISA/US

Signature of Authorized Officer

NGUYEN NGOC BO
INTERNATIONAL DIVISION
CHARLES ATKINSON

FURTHER INFORMATION CONTINUED FROM THE SECOND SHEET

- | | | |
|---|--|------|
| A | US, A, 4,467,421 (WHITE) 21 August 1984.
See the entire document. | 1-79 |
| A | IBM Technical Disclosure Bulletin 1983, Hempy,
"Multipathing in Peripheral Storage Systems",
see pages 3820-1. | 1-79 |
| A | US, A, 4,899,342 (POTTER ET AL.) 06 February
1990. See the entire document. | 1-79 |

V. OBSERVATIONS WHERE CERTAIN CLAIMS WERE FOUND UNSEARCHABLE¹

This international search report has not been established in respect of certain claims under Article 17(2) (a) for the following reasons:

1. Claim numbers because they relate to subject matter¹² not required to be searched by this Authority, namely:

2. Claim numbers because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out¹³, specifically:

3. Claim numbers because they are dependent claims not drafted in accordance with the second and third sentences of PCT Rule 6.4(a).

VI. OBSERVATIONS WHERE UNITY OF INVENTION IS LACKING¹⁴

This International Searching Authority found multiple inventions in this international application as follows:

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims of the international application.
2. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims of the international application for which fees were paid, specifically claims:
3. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claim numbers:
4. As all searchable claims could be searched without effort justifying an additional fee, the International Searching Authority did not invite payment of any additional fee.

Remark on Protest

- The additional search fees were accompanied by applicant's protest.
- No protest accompanied the payment of additional search fees.

THIS PAGE BLANK (USPTO)